arXiv:math/0506342v3 [math.ST] 22 Aug 2006

# Rodeo: Sparse Nonparametric Regression in High Dimensions

John Lafferty[1]
Computer Science Department
Machine Learning Department
Carnegie Mellon University

Larry Wasserman[2]
Department of Statistics
Machine Learning Department
Carnegie Mellon University

February 1, 2008

**Abstract**

We present a greedy method for simultaneously performing local bandwidth selection and variable selection in nonparametric regression. The method starts with a local linear estimator with large bandwidths, and incrementally decreases the bandwidth of variables for which the gradient of the estimator with respect to bandwidth is large. The method—called *rodeo* (regularization of derivative expectation operator)—conducts a sequence of hypothesis tests to threshold derivatives, and is easy to implement. Under certain assumptions on the regression function and sampling density, it is shown that the rodeo applied to local linear smoothing avoids the curse of dimensionality, achieving near optimal minimax rates of convergence in the number of relevant variables, as if these variables were isolated in advance.

**Keywords**: Nonparametric regression, sparsity, local linear smoothing, bandwidth estimation, variable selection, minimax rates of convergence.

[1]Research supported by NSF grants CCR-0122481, IIS-0312814, and IIS-0427206.
[2]Research supported by NIH grants R01-CA54852-07 and MH57881 and NSF grant DMS-0104016.

# I. Introduction

Estimating a high dimensional regression function is notoriously difficult due to the curse of dimensionality. Minimax theory precisely characterizes the curse. Let

$$Y_i = m(X_i) + \epsilon_i, \quad i = 1, \ldots, n \tag{1.1}$$

where $X_i = (X_i(1), \ldots, X_i(d)) \in \mathbb{R}^d$ is a $d$-dimensional covariate, $m : \mathbb{R}^d \to \mathbb{R}$ is the unknown function to estimate, and $\epsilon_i \sim N(0, \sigma^2)$. Then if $m$ is in $W_2(c)$, the $d$-dimensional Sobolev ball of order two and radius $c$, it is well known that

$$\liminf_{n\to\infty} n^{4/(4+d)} \inf_{\widehat{m}_n} \sup_{m \in W_2(c)} \mathcal{R}(\widehat{m}_n, m) \; > \; 0\,, \tag{1.2}$$

where $\mathcal{R}(\widehat{m}_n, m) = \mathbb{E}_m \int (\widehat{m}_n(x) - m(x))^2 \, dx$ is the risk of the estimate $\widehat{m}_n$ constructed from a sample of size $n$ (Györfi et al., 2002; Stone et al., 1997). Thus, the best rate of convergence is $n^{-4/(4+d)}$, which is impractically slow if $d$ is large.

However, for some applications it is reasonable to expect that the true function only depends on a small number of the total covariates. Suppose that $m$ satisfies such a sparseness condition, so that

$$m(x) = m(x_R) \tag{1.3}$$

where $x_R = (x_j : \; j \in R)$, $R \subset \{1, \ldots, d\}$ is a subset of the $d$ covariates, of size $r = |R| \ll d$. We call $\{x_j\}_{j \in R}$ the *relevant variables*. Note that if an oracle were to identify and isolate the relevant variables, the better minimax rate of $n^{-4/(4+r)}$ could be achieved, and this would be the fastest rate possible. Thus, we are faced with the problem of variable selection in nonparametric regression. Our strategy is to seek a greedy method that incrementally searches through bandwidths in small steps.

A large body of previous work has addressed this fundamental problem, which has led to a variety of methods to combat the curse of dimensionality. Many of these are based on very clever, though often heuristic techniques. For additive models of the form $m(x) = \sum_j m_j(x_j)$, standard methods like stepwise selection, $C_p$ and AIC can be used (Hastie et al., 2001). For spline models, Zhang et al. (2005) use likelihood basis pursuit, essentially the lasso adapted to the spline setting. CART (Breiman et al., 1984) and MARS (Friedman, 1991) effectively perform variable selection as part of their function fitting. Support vector regression can be seen as creating a sparse representation using basis pursuit in a reproducing kernel Hilbert space (Girosi, 1997). There is also a large literature on Bayesian methods, including methods for sparse Gaussian processes (Tipping, 2001; Smola and Bartlett, 2001; Lawrence et al.,

2003); see George and McCulloch (1997) for a brief survey. More recently, Li et al. (2005) use independence testing for variable selection and Bühlmann and Yu (2005) introduced a boosting approach. While these methods have met with varying degrees of empirical success, they can be challenging to implement and demanding computationally. Moreover, these methods are typically very difficult to analyze theoretically, and so come with limited formal guarantees. Indeed, the theoretical analysis of sparse *parametric* estimators such as the lasso (Tibshirani, 1996) is challenging, and only recently has significant progress been made on this front in the statistics and signal processing communities (Donoho, 2004; Fu and Knight, 2000; Tropp, 2004, 2006; Fan and Peng, 2004; Fan and Li, 2001).

In this paper we present a new approach for sparse nonparametric function estimation that is both computationally simple and amenable to theoretical analysis. We call the general framework *rodeo*, for "regularization of derivative expectation operator." It is based on the idea that bandwidth and variable selection can be simultaneously performed by computing the infinitesimal change in a nonparametric estimator as a function of the smoothing parameters, and then thresholding these derivatives to get a sparse estimate. As a simple version of this principle we use hard thresholding, effectively carrying out a sequence of hypothesis tests. A modified version that replaces testing with soft thresholding may be viewed as solving a sequence of lasso problems. The potential appeal of this approach is that it can be based on relatively simple and theoretically well understood nonparametric techniques such as local linear smoothing, leading to methods that are simple to implement and can be used in high dimensional problems. Moreover, we show that they can achieve near optimal minimax rates of convergence, and therefore circumvent the curse of dimensionality when the true function is indeed sparse. When applied in one dimension, our method yields a local bandwidth selector and is similar to the estimators of Ruppert (1997) and Lepski et al. (1997). The method in Lepski et al. (1997) and its multivariate extension in Kerkyacharian, Lepski and Picard (2001) yield estimators that are more refined than our method in the sense that their estimator is spatially adaptive over large classes of function spaces. However, their method is not greedy: it involves searching over a large class of bandwidths. Our goal is to develop a greedy method that scales to high dimensions.

Our method is related to the structural adapation method of Hristache et al. (2001) and Samarov et al. (2005), which is designed for multi-index models. The general multi-index model is

$$Y = g_0(Tx) + \epsilon \tag{1.4}$$

where $x \in \mathbb{R}^d$ and $T$ is a linear orthonormal mapping from $\mathbb{R}^d$ onto $\mathbb{R}^r$ with $r < d$. Variable

selection corresponds to taking $T$ to be a $r$ by $d$ matrix of 0's and 1's with each $T_{ij} = 1$ if $x_j$ is the $i^{\text{th}}$ relevant variable. Nonparametric variable selection can also be regarded as a special case of the partially linear model in Samarov et al. (2005), which takes

$$Y = \theta^T x_1 + G(x_2) + \epsilon \tag{1.5}$$

where $x = (x_1, x_2)$. Taking $\theta$ to be zero yields the model in this paper. The advantage of structural adaptation is that it yields, under certain conditions, $\sqrt{n}$ estimates of the image of $T$ in (1.4) and $\theta$ in (1.5). However, structural adaptation does not yield optimal bandwidths or optimal estimates of the regression function, although this is not the intended goal of the method.

In the following section we outline the basic rodeo approach, which is actually a general strategy that can be applied to a wide range of nonparametric estimators. We then specialize in Section 3 to the case of local linear smoothing, since the asymptotic properties of this smoothing technique are fairly well understood. In particular, we build upon the analysis of Ruppert and Wand (1994) for local linear regression; a notable difference is that we allow the dimension to increase with sample size, which requires a more detailed analysis of the asymptotics. In Section 4 we present some simple examples of the rodeo, before proceeding to an analysis of its properties in Section 5. Our main theoretical result characterizes the asymptotic running time, selected bandwidths, and risk of the algorithm. Finally, in Section 6 we present further examples and discuss several extensions of the basic version of the rodeo considered in the earlier sections. The proofs of the theoretical properties of the rodeo are given in Section 7.

## II. Rodeo: The Main Idea

The key idea in our approach is as follows. Fix a point $x$ and let $\widehat{m}_h(x)$ denote an estimator of $m(x)$ based on a vector of smoothing parameters $h = (h_1, \ldots, h_d)$. If $c$ is a scalar, then we write $h = c$ to mean $h = (c, \ldots, c)$.

Let $M(h) = \mathbb{E}(\widehat{m}_h(x))$ denote the mean of $\widehat{m}_h(x)$. For now, assume that $x = x_i$ is one of the observed data points and that $\widehat{m}_0(x) = Y_i$. In that case, $m(x) = M(0) = \mathbb{E}(Y_i)$. If $P = (h(t) : \ 0 \leq t \leq 1)$ is a smooth path through the set of smoothing parameters with

$h(0) = 0$ and $h(1) = 1$ (or any other fixed, large bandwidth) then

$$\begin{aligned} m(x) &= M(0) = M(1) + M(0) - M(1) && \text{(2.1a)} \\ &= M(1) - \int_0^1 \frac{dM(h(s))}{ds}\, ds && \text{(2.1b)} \\ &= M(1) - \int_0^1 \big\langle D(s), \dot{h}(s) \big\rangle ds && \text{(2.1c)} \end{aligned}$$

where

$$D(h) = \nabla M(h) = \left( \frac{\partial M}{\partial h_1}, \ldots, \frac{\partial M}{\partial h_d} \right)^T \tag{2.2}$$

is the gradient of $M(h)$ and $\dot{h}(s) = \frac{dh(s)}{ds}$ is the derivative of $h(s)$ along the path. A biased, low variance estimator of $m(x)$ is $\widehat{m}_1(x)$. An unbiased estimator of $D(h)$ is

$$Z(h) = \left( \frac{\partial \widehat{m}_h(x)}{\partial h_1}, \ldots, \frac{\partial \widehat{m}_h(x)}{\partial h_d} \right)^T . \tag{2.3}$$

The naive estimator

$$\widehat{m}(x) = \widehat{m}_1(x) - \int_0^1 \big\langle Z(s), \dot{h}(s) \big\rangle ds \tag{2.4}$$

is identically equal to $\widehat{m}_0(x) = Y_i$, which has poor risk since the variance of $Z(h)$ is large for small $h$. However, our sparsity assumption on $m$ suggests that there should be paths for which $D(h)$ is also sparse. Along such a path, we replace $Z(h)$ with an estimator $\widehat{D}(h)$ that makes use of the sparsity assumption. Our estimate of $m(x)$ is then

$$\widetilde{m}(x) = \widehat{m}_1(x) - \int_0^1 \big\langle \widehat{D}(s), \dot{h}(s) \big\rangle ds\,. \tag{2.5}$$

To implement this idea we need to do two things: (i) we need to find a path for which the derivative is sparse and (ii) we need to take advantage of this sparseness when estimating $D$ along that path.

The key observation is that if $x_j$ is irrelevant, then we expect that changing the bandwidth $h_j$ for that variable should cause only a small change in the estimator $\widehat{m}_h(x)$. Conversely, if $x_j$ is relevant, then we expect that changing the bandwidth $h_j$ for that variable should cause a large change in the estimator. Thus, $Z_j = \partial \widehat{m}_h(x)/\partial h_j$ should discriminate between relevant and irrelevant covariates. To simplify the procedure, we can replace the continuum of bandwidths with a discrete set where each $h_j \in \mathcal{B} = \{h_0, \beta h_0, \beta^2 h_0, \ldots\}$ for some $0 < \beta < 1$. Moreover, we can proceed in a greedy fashion by estimating $D(h)$ sequentially with $h_j \in \mathcal{B}$ and setting $\widehat{D}_j(h) = 0$ when $h_j < \widehat{h}_j$, where $\widehat{h}_j$ is the first $h$ such that $|Z_j(h)| < \lambda_j(h)$ for

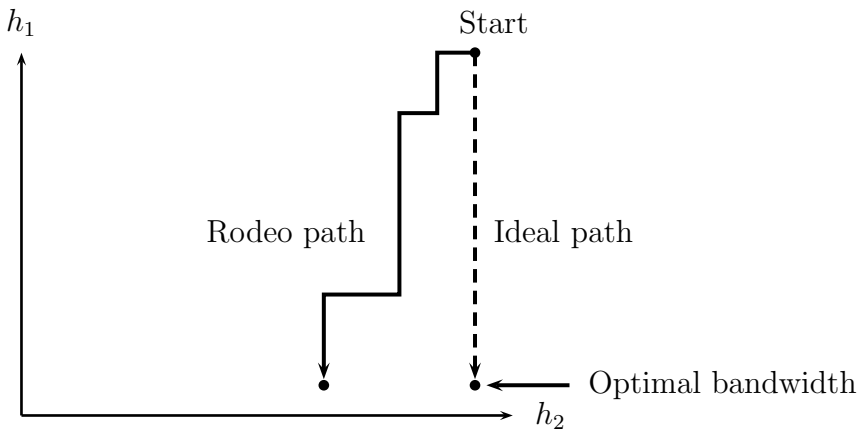


Figure 1: Conceptual illustration: The bandwidths for the relevant variables ($h_1$) are shrunk, while the bandwidths for the irrelevant variables ($h_2$) are kept relatively large.

some threshold $\lambda_j$. This greedy version, coupled with the hard threshold estimator, yields $\widetilde{m}(x) = \widehat{m}_{\widehat{h}}(x)$. A conceptual illustration of the idea is shown in Figure 1.

To further elucidate the idea, consider now the one-dimensional case $x \in \mathbb{R}$, so that

$$m(x) = M(1) - \int_0^1 \frac{dM(h)}{dh} dh = M(1) - \int_0^1 D(h)\, dh. \tag{2.6}$$

Suppose that $\widehat{m}_h(x) = \sum_{i=1}^n Y_i\, \ell_i(x,h)$ is a linear estimator, where the weights $\ell_i(x,h)$ depend on a bandwidth $h$.

In this case

$$Z(h) = \sum_{i=1}^n Y_i\, \ell_i'(x,h) \tag{2.7}$$

where the prime denotes differentiation with respect to $h$. Then we set

$$\widetilde{m}(x) = \widehat{m}_1(x) - \int_0^1 \widehat{D}(h)\, dh \tag{2.8}$$

where $\widehat{D}(h)$ is an estimator of $D(h)$. Now,

$$Z(h) \approx N(b(h), s^2(h)) \tag{2.9}$$

where, for typical smoothers, $b(h) \approx Ah$ and $s^2(h) \approx C/nh^3$ for some constants $A$ and $C$. Take the hard threshold estimator

$$\widehat{D}(h) = Z(h) I\big(|Z(h)| > \lambda(h)\big) \tag{2.10}$$

where $\lambda(h)$ is chosen to be slightly larger than $s(h)$. An alternative is the soft-threshold estimator

$$\widehat{D}(h) = \text{sign}(Z(h))(|Z(h)| - \lambda(h))_+. \tag{2.11}$$

The greedy algorithm, coupled with the hard threshold estimator, yields a bandwidth selection procedure based on testing. This approach to bandwidth selection is very similar to that of Lepski et al. (1997), who take

$$\widehat{h} = \max\{h \in \mathcal{H} : \ \phi(h, \eta) = 0 \text{ for all } \eta < h\} \tag{2.12}$$

where $\phi(h, \eta)$ is a test for whether $\widehat{m}_\eta$ improves on $\widehat{m}_h$. This more refined test leads to estimators that achieve good spatial adaptation over large function classes. Kerkyacharian, Lepski and Picard (2001) extend the idea to multiple dimensions. Our approach is also similar to a method of Ruppert (1997) that uses a sequence of decreasing bandwidths and then estimates the optimal bandwidth by estimating the mean squared error as a function of bandwidth. Our greedy approach only tests whether an infinitesimal change in the bandwidth from its current setting leads to a significant change in the estimate, and is more easily extended to a practical method in higher dimensions.

## III. Rodeo Using Local Linear Regression

Now we present the multivariate rodeo in detail. We use local linear smoothing as the basic method since it is known to have many good properties. Let $x = (x(1), \ldots, x(d))$ be some target point at which we want to estimate $m$. Let $\widehat{m}_H(x)$ denote the local linear estimator of $m(x)$ using bandwidth matrix $H$. Thus,

$$\widehat{m}_H(x) = e_1^T (X_x^T W_x X_x)^{-1} X_x^T W_x Y \equiv S_x Y \tag{3.1}$$

where $e_1 = (1, 0, \ldots, 0)^T$,

$$X_x = \begin{pmatrix} 1 & (X_1 - x)^T \\ \vdots & \vdots \\ 1 & (X_n - x)^T \end{pmatrix}, \tag{3.2}$$

$W_x$ is diagonal with $(i, i)$ element $K_H(X_i - x)$ and $K_H(u) = |H|^{-1/2} K(H^{-1/2} u)$. The estimator $\widehat{m}_H$ can be written as

$$\widehat{m}_H(x) = \sum_{i=1}^{n} G(X_i, x, h)\, Y_i \tag{3.3}$$

where

$$G(u,x,h) = e_1^T (X_x^T W_x X_x)^{-1} \begin{pmatrix} 1 \\ (u-x)^T \end{pmatrix} K_H(u-x) \tag{3.4}$$

is called the *effective kernel*. One can regard local linear regression as a refinement of kernel regression where the effective kernel $G$ adjusts for boundary bias and design bias; see Fan (1992), Hastie and Loader (1993) and Ruppert and Wand (1994).

We assume that the covariates are random with density $f(x)$ and that $x$ is interior to the support of $f$. We make the same assumptions as Ruppert and Wand (1994) in their analysis of the bias and variance of local linear regression. In particular:

(i) The kernel $K$ has compact support with zero odd moments and there exists $\nu_2 = \nu_2(K) \neq 0$ such that

$$\int uu^\top K(u)\, du = \nu_2(K) I \tag{3.5}$$

where $I$ is the $d \times d$ identity matrix.

(ii) The sampling density $f(x)$ is continuously differentiable and strictly positive.

In the version of the algorithm that follows, we take $K$ to be a product kernel and $H$ to be diagonal with elements $h = (h_1, \ldots, h_d)$ and we write $\widehat{m}_h$ instead of $\widehat{m}_H$.

Our method is based on the statistic

$$Z_j = \frac{\partial \widehat{m}_h(x)}{\partial h_j} = \sum_{i=1}^n G_j(X_i, x, h) Y_i \tag{3.6}$$

where

$$G_j(u,x,h) = \frac{\partial G(u,x,h)}{\partial h_j}. \tag{3.7}$$

Let

$$\mu_j \equiv \mu_j(h) = \mathbb{E}(Z_j | X_1, \ldots, X_n) = \sum_{i=1}^n G_j(X_i, x, h) m(X_i) \tag{3.8}$$

and

$$s_j^2 \equiv s_j^2(h) = \mathbb{V}(Z_j | X_1, \ldots, X_n) = \sigma^2 \sum_{i=1}^n G_j(X_i, x, h)^2. \tag{3.9}$$

In Section 4.C we explain how to estimate $\sigma$; for now, assume that $\sigma$ is known. The hard thresholding version of the rodeo algorithm is described in Figure 2. We make use of a sequence $c_n$ satisfying $dc_n = \Omega(\log n)$, where we write $f(n) = \Omega(g(n))$ if $\liminf_n \left| \frac{f(n)}{g(n)} \right| > 0$.

*Rodeo: Hard thresholding version*

1. *Select* $\beta_n = n^{-\alpha/\log^3 n}$ for some $0 < \alpha < 1$ and initial bandwidth

$$h_0 = \frac{c_0}{\log\log n} \tag{3.10}$$

   for some constant $c_0$. Let $c_n$ be a sequence satisfying $c_n = O(1)$.

2. *Initialize* the bandwidths, and activate all covariates:

   (a) $h_j = h_0$, $j = 1, 2, \ldots, d$.

   (b) $\mathcal{A} = \{1, 2, \ldots, d\}$

3. *While* $\mathcal{A}$ *is nonempty*, do for each $j \in \mathcal{A}$:

   (a) Compute the estimated derivative expectation: $Z_j$ (equation 3.6) and $s_j$ (equation 3.9).

   (b) Compute the threshold $\lambda_j = s_j\sqrt{2\log(nc_n)}$.

   (c) If $|Z_j| > \lambda_j$, then set $h_j \leftarrow \beta h_j$; otherwise remove $j$ from $\mathcal{A}$.

4. *Output* bandwidths $h^\star = (h_1, \ldots, h_d)$ and estimator $\widetilde{m}(x) = \widehat{m}_{h^\star}(x)$.

Figure 2: The hard thresholding version of the rodeo, which can be applied using the derivatives $Z_j$ of any nonparametric smoother.

To derive an explicit expression for $Z_j$, equivalently $G_j$, we use

$$\frac{\partial A^{-1}}{\partial h} = -A^{-1}\frac{\partial A}{\partial h}A^{-1} \tag{3.11}$$

to get that

$$Z_j = \frac{\partial \widehat{m}_h(x)}{\partial h_j} \tag{3.12a}$$

$$= e_1^\top (X^\top W X)^{-1} X^\top \frac{\partial W}{\partial h_j} Y - e_1^\top (X^\top W X)^{-1} X^\top \frac{\partial W}{\partial h_j} X (X^\top W X)^{-1} X^\top W Y \tag{3.12b}$$

$$= e_1^\top (X^\top W X)^{-1} X^\top \frac{\partial W}{\partial h_j} (Y - X\widehat{\alpha}) \tag{3.12c}$$

where $\widehat{\alpha} = (X^\top W X)^{-1} X^\top W Y$ is the coefficient vector for the local linear fit (and we have dropped the dependence on the local point $x$ in the notation).

Note that the factor $|H|^{-1} = \prod_{i=1}^d 1/h_i$ in the kernel cancels in the expression for $\widehat{m}$, and therefore we can ignore it in our calculation of $Z_j$. Assuming a product kernel we have

$$W = \text{diag}\left(\prod_{j=1}^d K((X_{1j} - x_j)/h_j), \ldots, \prod_{j=1}^d K((X_{nj} - x_j)/h_j)\right) \tag{3.13}$$

and $\partial W / \partial h_j = W L_j$ where

$$L_j = \text{diag}\left(\frac{\partial \log K((X_{1j} - x_j)/h_j)}{\partial h_j}, \ldots, \frac{\partial \log K((X_{nj} - x_j)/h_j)}{\partial h_j}\right) \tag{3.14}$$

and thus

$$Z_j = e_1^\top (X^\top W X)^{-1} X^\top W L_j (Y - X\widehat{\alpha}) = e_1^\top B L_j (I - XB) Y = G_j(x,h)^\top Y \tag{3.15}$$

where $B = (X^\top W X)^{-1} X^\top W$.

The calculation of $L_j$ is typically straightforward. As two examples, with the Gaussian kernel $K(u) = \exp(-u^2/2)$ we have

$$L_j = \frac{1}{h_j^3} \text{diag}\left((X_{1j} - x_j)^2, \ldots, (X_{nj} - x_j)^2\right) \tag{3.16}$$

and for the Epanechnikov kernel $K(u) = (5 - x^2)\, \mathbb{I}(|x| \leq \sqrt{5})$ we have

$$L_j = \frac{1}{h_j^3} \text{diag}\left(\frac{2(X_{1j} - x_j)^2}{5 - (X_{1j} - x_j)^2/h_j^2}\, \mathbb{I}(|X_{1j} - x_j| \leq \sqrt{5} h_j), \ldots, \right. \tag{3.17a}$$

$$\left. \frac{2(X_{nj} - x_j)^2}{5 - (X_{nj} - x_j)^2/h_j^2}\, \mathbb{I}(|X_{1j} - x_j| \leq \sqrt{5} h_j)\right) \tag{3.17b}$$

## IV. Examples

In this section we illustrate the rodeo on some examples. We return to the examples later when we discuss estimating $\sigma$, as well as a global (non-local) version of the rodeo.

### *A. Two Relevant Variables*

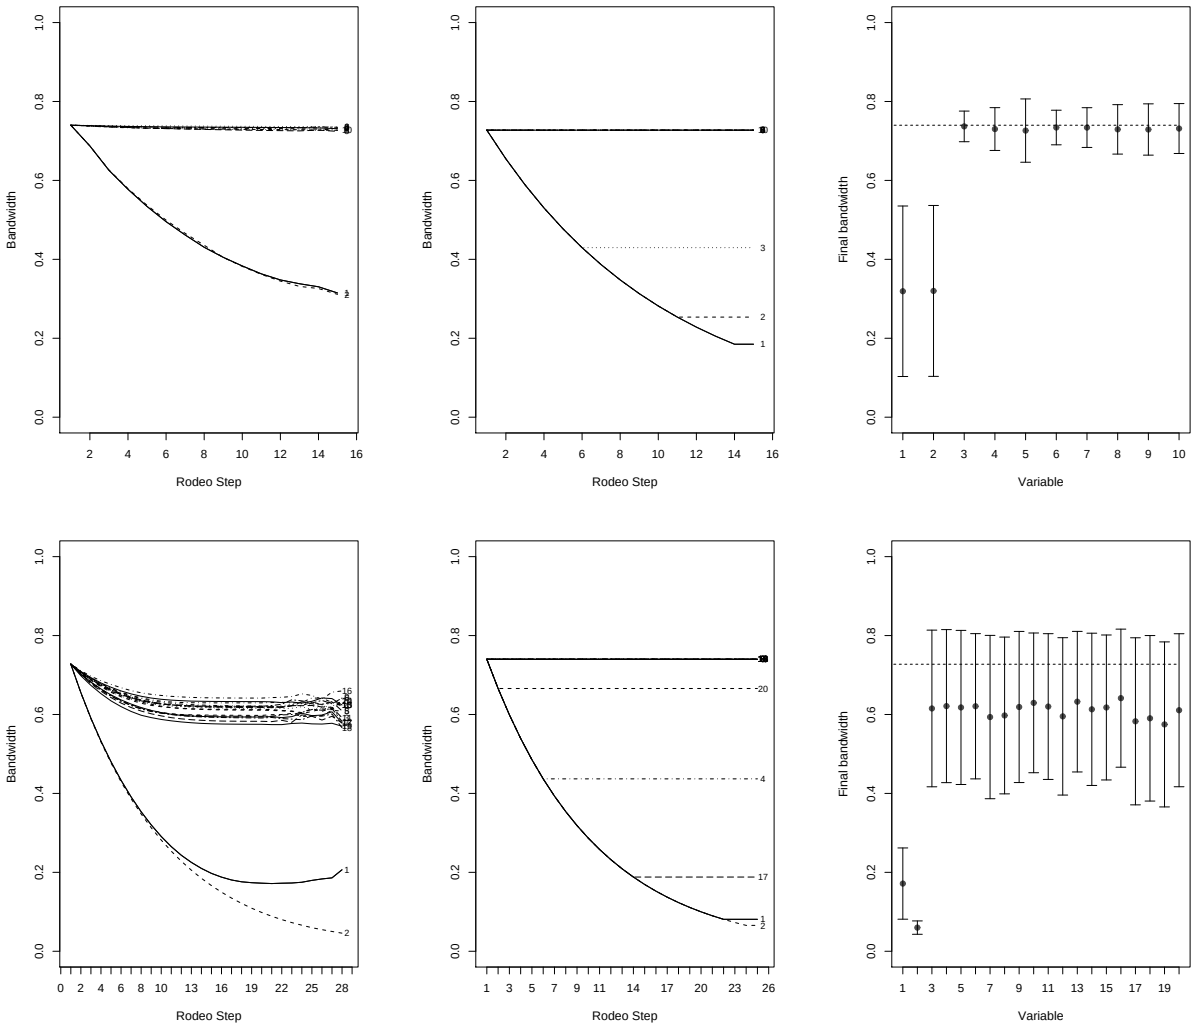


Figure 3: Rodeo run on synthetic data sets, showing average bandwidths over 200 runs (left), final bandwidths with standard errors (right), and bandwidths on a single run of the algorithm (center). In the top plots the regression function is $m(x) = 5x_1^2x_2^2$ with $d = 10$, $n = 500$, and $\sigma = .5$ and in the lower plots the regression function is $m(x) = 2(x_1 + 1)^3 + 2\sin(10x_2)$, $d = 20$, $n = 750$, and $\sigma = 1$. The figures show that the bandwidths for the relevant variables $x_1$ and $x_2$ are shrunk, while the bandwidths for the irrelevant variables remain large.

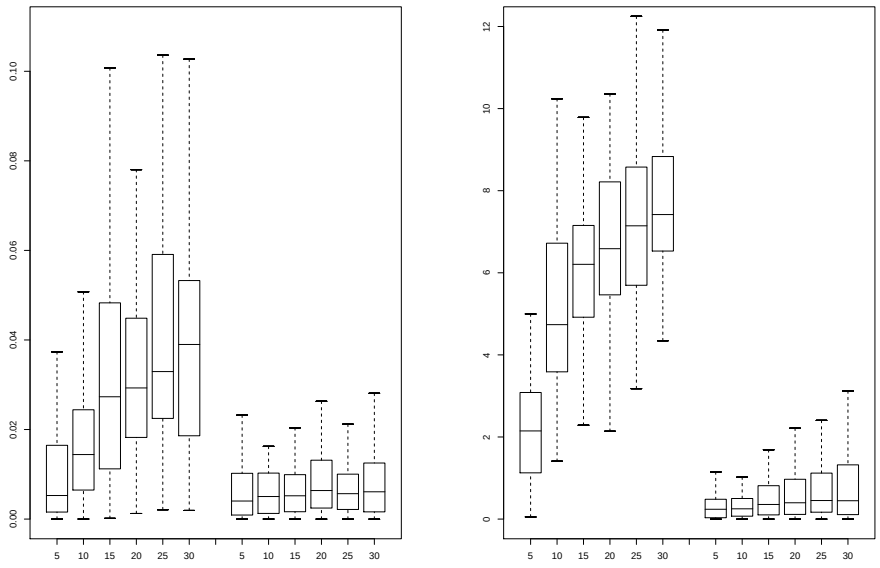


Figure 4: Squared error of the estimator on the previous examples, $m(x) = 5x_1^2x_2^2$ (left) and $m(x) = 2(x_1 + 1)^3 + 2\sin(10x_2)$ right. For each plot, the left six boxplots show the risk in different dimensions ($d = 5, 10, 15, 20, 25, 30$) when using a single bandwidth, chosen by leave-one-out cross validation. The right six boxplots show the squared error on the same data with bandwidths selected using the rodeo.

In the first example, we take $m(x) = 5x_1^2x_2^2$ with $d = 10$, $\sigma = .5$ with $x_i \sim \text{Uniform}(0, 1)$. The algorithm is applied to the local linear estimates around the test point $x_0 = (\frac{1}{2}, \ldots, \frac{1}{2})$, with $\beta = 0.8$. Figure 3 shows the bandwidths averaged over 200 runs of the rodeo, on data sets of size $n = 750$. The second example in Figure 3 shows the algorithm applied to the function $m(x) = 2(x_1 + 1)^3 + 2\sin(10x_2)$, in this case in $d = 20$ dimensions with $\sigma = 1$.

The plots demonstrate how the bandwidths $h_1$ and $h_2$ of the relevant variables are shrunk, while the bandwidths of the relevant variables tend to remain large.

*B. A One-Dimensional Example*

The next figure illustrates the algorithm in one dimension. The underlying function in this case is $m(x) = (1/x)\sin(15/x)$, and $n = 1500$ data points are sampled as $x \sim \text{Uniform}(0, 1) + \frac{1}{2}$. The algorithm is run at two test points; the function is more rapidly varying near the test point $x = 0.67$ than near the test point $x = 1.3$, and the rodeo appropriately selects a smaller

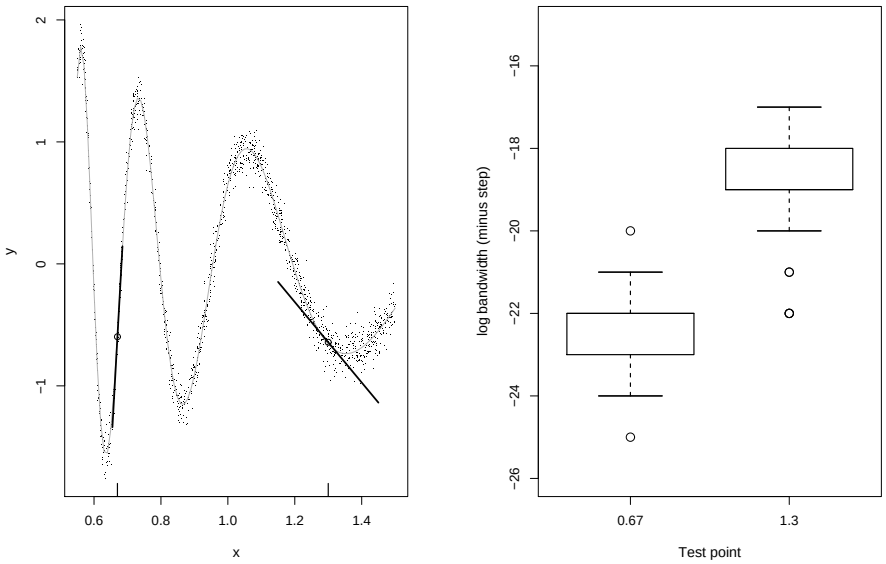


Figure 5: A one-dimensional example. The regression function is $m(x) = (1/x)\sin(15/x)$, and $n = 1500$ data points are sampled, $x \sim \text{Uniform}(0,1) + \frac{1}{2}$. The left plot shows the local linear fit at two test points; the right plot shows the final log bandwidth, $\log_{1/\beta} h_\star$, (equivalently, minus the number of steps) of the rodeo over 50 randomly generated data sets.

bandwidth at $x = 0.67$. The right plot of Figure 5 displays boxplots for logarithm of the final bandwidth, in the base $1/\beta$ (equivalently, minus the number of steps in the algorithm), averaged over 50 randomly generated data sets.

The figure illustrates how smaller bandwidths are selected where the function is more rapidly varying. However, we do not claim that the method is adaptive over large classes of function spaces. As discussed earlier, the technique is intentionally a greedy algorithm; adapting to unknown smoothness may require a more refined search over bandwidths that does not scale to large dimensions, and is out of the scope of the current paper.

*C. Estimating $\sigma$*

The algorithm requires that we insert an estimate $\widehat{\sigma}$ of $\sigma$ in (3.9). An estimator for $\sigma$ can be obtained by generalizing a method of Rice (1984). For $i < \ell$, let

$$d_{i\ell} = \|X_i - X_\ell\| \,. \tag{4.1}$$

Fix an integer $J$ and let $\mathcal{E}$ denote the set of pairs $(i, \ell)$ corresponding the $J$ smallest values of $d_{i\ell}$. Now define

$$\widehat{\sigma}^2 = \frac{1}{2J} \sum_{i,\ell \in \mathcal{E}} (Y_i - Y_\ell)^2. \tag{4.2}$$

Then,

$$\mathbb{E}(\widehat{\sigma}^2) \;=\; \sigma^2 + \text{bias} \tag{4.3}$$

where

$$\text{bias} \;\leq\; D \sup_x \sum_{j=1}^{r} \left| \frac{\partial m_j(x)}{\partial x_j} \right| \tag{4.4}$$

with $D$ given by

$$D = \max_{i,\ell \in \mathcal{E}} \|X_i - X_\ell\|. \tag{4.5}$$

There is a bias-variance tradeoff: large $J$ makes $\widehat{\sigma}^2$ positively biased, and small $J$ makes $\widehat{\sigma}^2$ highly variable. Note, however, that the bias is mitigated by sparsity (small $r$).

A more robust estimate may result from taking

$$\widehat{\sigma}^2 = \frac{\sqrt{\pi}}{2} \,\text{median}\, \{|Y_i - Y_\ell|\}_{i,\ell \in \mathcal{E}} \tag{4.6}$$

where the constant comes from observing that if $X_i$ is close to $X_\ell$, then

$$|Y_i - Y_\ell| \sim |N(0, 2\sigma^2)| = \sqrt{2}\sigma |Z|, \tag{4.7}$$

where $Z$ is a standard normal with $\mathbb{E}|Z| = \sqrt{2/\pi}$.

Now we redo the earlier examples, taking $\sigma$ as unknown. Figure 6 shows the result of running the algorithm on the examples of Section 4.A, however now estimating the noise using estimate (4.6). For the higher dimensional example, with $d = 20$, the noise variance is over-estimated, with the primary result that the irrelevant variables are more aggressively thresholded out; compare Figure 6 to Figure 3.

Although we do not pursue it in this paper, there is also the possibility of allowing $\sigma(x)$ to be a function of $x$ and estimating it locally.

### *D. Computational Cost*

When based on a local linear estimator, each step of the rodeo algorithm has the same computational cost as constructing a single local linear fit. This is dominated by the cost of constructing the matrix inverse $(X^T W X)^{-1}$ in equation (3.15). Since the derivative needs

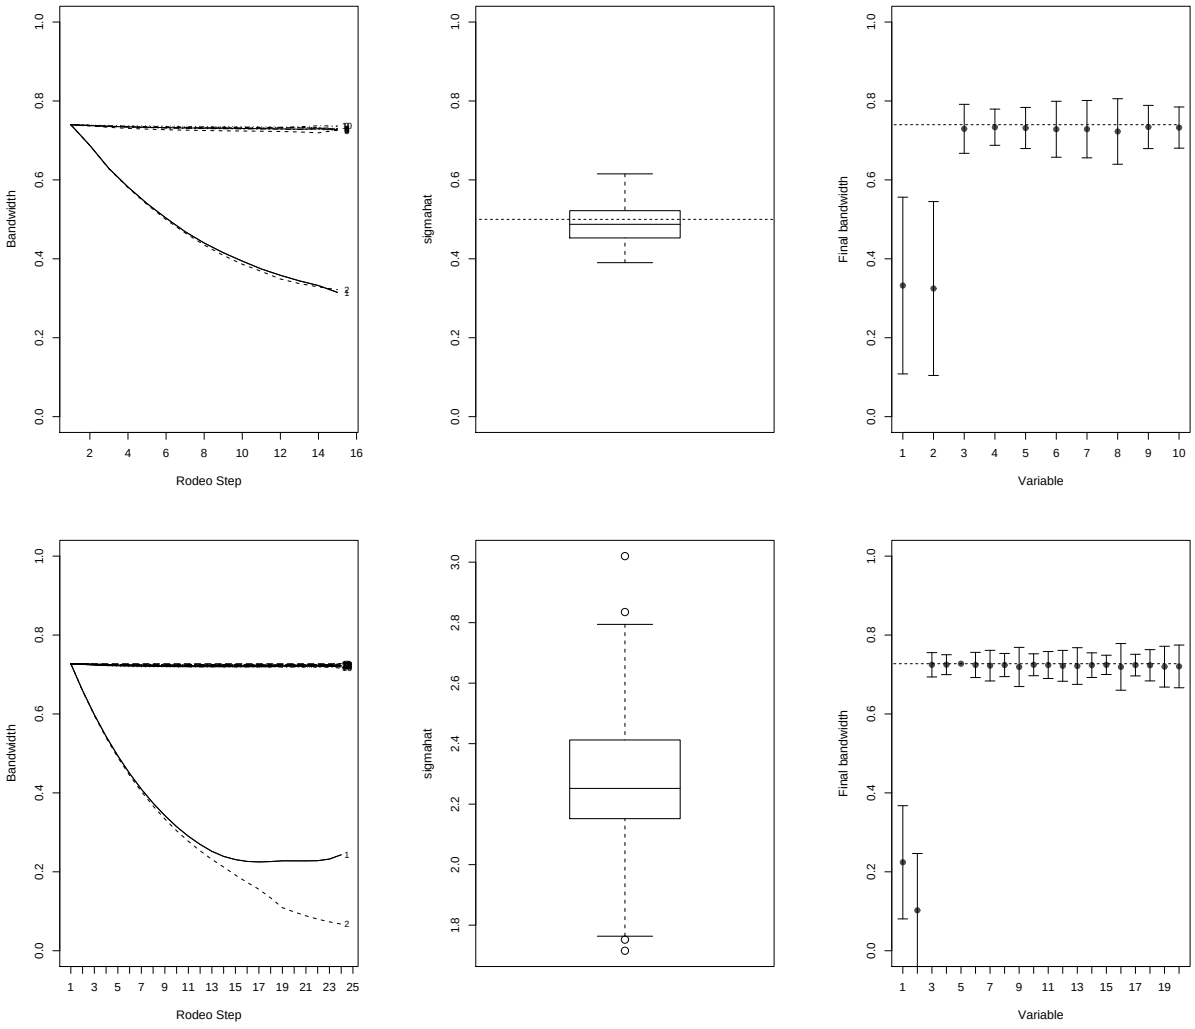


Figure 6: Rodeo run on the examples of Section 4.A, but now estimating the noise using the estimate $\widehat{\sigma}$ discussed in Section 4.C. Top: $\sigma = .5$, $d = 10$; bottom: $\sigma = 1$, $d = 20$. In higher dimensions the noise is over-estimated (center plots), which results in the irrelevant variables being more aggressively eliminated; compare Figure 3.

to be computed for every variable, the algorithm thus scales as $O(d^4)$ in the dimension $d$. Implemented in R, the 20 dimensional example in Figure 3 takes 4 hours, 4 minutes and 40 seconds for 200 runs, or 73.4 seconds per run, when executed on a 1.5 GHz PowerPC Macintosh laptop. Although we focus on local linear regression, it should be noted that very similar results are obtained with kernel regression, which requires no matrix inversion. Using kernel regression, the same example requires 12 minutes and 33 seconds, or 3.7 seconds per run.

## V. Properties of the Rodeo

We now give some results on the statistical properties of the hard thresholding version of the rodeo estimator. Formally, we use a triangular array approach so that $m(x)$, $f(x)$, $d$ and $r$ can all change as $n$ changes, although we often suppress the dependence on $n$ for notational clarity. We assume throughout that $m$ has continuous third order derivatives in a neighborhood of $x$. For convenience of notation we assume that the covariates are numbered such that the relevant variables $x_j$ correspond to $1 \leq j \leq r$ and the irrelevant variables $x_j$ correspond to $r+1 \leq j \leq d$.

A key aspect of our analysis is that we allow the dimension $d$ to increase with sample size $n$, and show that the algorithm achieves near optimal minimax rates of convergence if $d = O(\log n / \log\log n)$. This hinges on a careful analysis of the asymptotic bias and variance of the estimated derivative $Z_j$, taking the increasing dimension into account. We conjecture that, without further assumptions, $d$ cannot increase at a significantly faster rate, while obtaining near optimal rates of convergence.

The results are stated below, with the complete proofs given in Section 7.

We write $Y_n = \widetilde{O}_P(a_n)$ to mean that $Y_n = O_P(b_n a_n)$ where $b_n$ is logarithmic in $n$. As noted earlier, we write $a_n = \Omega(b_n)$ if $\liminf_n \left|\frac{a_n}{b_n}\right| > 0$; similarly $a_n = \widetilde{\Omega}(b_n)$ if $a_n = \Omega(b_n c_n)$ where $c_n$ is logarithmic in $n$.

Let $\mathcal{H} = \begin{pmatrix} \mathcal{H}_R & 0 \\ 0 & 0 \end{pmatrix}$ denote the Hessian of $m(x)$, let $h_j^{(t)}$ denote the $j^{\text{th}}$ bandwidth at step $t$ of the algorithm and denote the bandwidth matrix by $H^{(t)} = \text{diag}((h_1^{(t)})^2, \ldots, (h_d^{(t)})^2)$.

Similarly, let $H_R^{(t)} = \text{diag}((h_1^{(t)})^2, \ldots, (h_r^{(t)})^2)$. Define

$$\mu_j^{(t)} = \frac{\partial}{\partial h_j} \mathbb{E}\left[\widehat{m}_{H^{(t)}}(x) - m(x) \,|\, X_1, \ldots, X_n\right]. \tag{5.1}$$

Our first result is an asymptotic expansion of the bias of the estimator.

**Theorem 5.1.** *Suppose that $x$ is interior to the support of $f$ and that $K$ is a product kernel. Assume that*

$$d = O(\log n / \log\log n), \qquad r = O(1), \tag{5.2}$$

$$h_0 = \frac{c_0}{\log\log n} \tag{5.3}$$

*for some $c_0 > 0$ and*

$$\beta_n = n^{-\alpha/\log^3 n} \tag{5.4}$$

*for some $0 < \alpha < 1$. Define*

$$L_j^{(t)} = \begin{cases} \nu_2 m_{jj}(x) h_j^{(t)} & j \le r \\ -\,tr\left(H_R^{(t)} \mathcal{H}_R^{(t)}\right) \nu_2^2 \left(\nabla_j \log f(x)\right)^2 h_j^{(t)} & j > r. \end{cases} \tag{5.5}$$

*where $\nu_2$ is defined in equation* (3.5)*. Suppose that*

$$\max_{a,b} \sup_x \left| \frac{\partial^2 f(x)}{\partial x_a \partial x_b} \right| = O(1/d^2) \tag{5.6}$$

*and*

$$\max_{a,b,c} \sup_x \left| \frac{\partial^3 f(x)}{\partial x_a \partial x_b \partial x_c} \right| = O(1/d^3). \tag{5.7}$$

*Then, for $T_n \le c_1 \log n$,*

$$\max_{\substack{1 \le s \le T_n \\ 1 \le j \le d}} |\mu_j^{(t)} - L_j^{(t)}| = O_P\left((h_j^{(t)})^2 / d\right). \tag{5.8}$$

The next result characterizes the variance $v_j^{(t)} = \text{Var}(Z_j(t) \,|\, X_1, \ldots, X_n)$.

**Theorem 5.2.** *Assume the conditions of Theorem 5.1. Let*

$$s_j^2(t) = \frac{C}{n(h_j^{(t)})^2} \prod_{k=1}^{d} \frac{1}{h_k^{(t)}} \tag{5.9}$$

*where $h_j^{(t)}$ is the $j^{\text{th}}$ bandwidth at step $t$, with $C = \sigma^2 \int K^2(u)\,du/f(x)$. Then,*

$$\mathbb{P}\left( \max_{\substack{1 \le s \le T_n \\ 1 \le j \le d}} \left| \frac{v_j^{(t)}}{s_j^2(t)} - 1 \right| > \epsilon \right) \to 0, \tag{5.10}$$

*for all $\epsilon > 0$.*

Our main theoretical result characterizes the asymptotic running time, selected bandwidths, and risk of the algorithm. In order to get a practical algorithm, we need to make assumptions on the functions $m$ and $f$.

(A1) For each $j > r$,

$$\nabla_j \log f(x) = O\left( \frac{1}{n^{1/4}} \right). \tag{5.11}$$

(A2)

$$\liminf_n \min_{1 \le j \le r} |m_{jj}(x)| > 0\,. \tag{5.12}$$

*Explanation of the Assumptions.* To give the intuition behind these assumptions, recall from Theorem 5.1 that

$$\mu_j = \begin{cases} A_j h_j + o_P(h_j) & j \le r \\ B_j h_j + o_P(h_j) & j > r \end{cases} \tag{5.13}$$

where

$$A_j = \nu_2 m_{jj}(x), \qquad B_j = -\operatorname{tr}(H_R \mathcal{H}_R)\nu_2^2 (\nabla_j \log f(x))^2. \tag{5.14}$$

Moreover, $\mu_j \approx 0$ when the sampling density $f$ is uniform or the data are on a regular grid. Consider assumption (A1). If $f$ is uniform then this assumption is automatically satisfied since then $\mu_j(s) \approx 0$ for $j > r$. More generally, $\mu_j$ is approximately proportional to $(\nabla_j \log f(x))^2$ for $j > r$ which implies that $|\mu_j| \approx 0$ for irrelevant variables if $f$ is sufficiently smooth in the variable $x_j$. Hence, assumption (A1) can be interpreted as requiring that $f$ is sufficiently smooth in the irrelevant dimensions. If $d$ is constant and $f$ is smooth, then this assumption is *implied* by Theorem 5.1. However, we allow for the possibility that $d$ grows with $n$.

Now consider assumption (A2). Equation (5.13) ensures that $\mu_j$ is proportional to $h_j|m_{jj}(x)|$ for small $h_j$. Since we take the initial bandwidth $h_0$ to be decreasingly slowly with $n$, (A2) implies that $|\mu_j(h)| \ge c h_j |m_{jj}(x)|$ for some constant $c > 0$, for sufficiently large $n$.

**Theorem 5.3.** *Assume the conditions of Theorem 5.1 and suppose that assumptions (A1) and (A2) hold. In addition, suppose that* $A_{min} = \min_{j \le r} |m_{jj}(x)| = \widetilde{\Omega}(1)$ *and* $A_{max} = \max_{j \le r} |m_{jj}(x)| = \widetilde{O}(1)$. *Then the rodeo outputs bandwidths* $h^\star$ *that satisfy*

$$\mathbb{P}\left(h_j^\star = h_0 \text{ for all } j > r\right) \longrightarrow 1 \tag{5.15}$$

*and for every* $\epsilon > 0$,

$$\mathbb{P}\left(n^{-1/(4+r)-\epsilon} \le h_j^\star \le n^{-1/(4+r)+\epsilon} \text{ for all } j \le r\right) \longrightarrow 1\,. \tag{5.16}$$

**Corollary 5.4.** *Under the conditions of Theorem 5.3, the risk of the rodeo estimator satisfies*

$$\mathcal{R}(h^\star) = \mathbb{E}\int (\widehat{m}_{h^\star}(x) - m(x))^2\, dx = O_P\left(n^{-4/(4+r)+\epsilon}\right) \tag{5.17}$$

*for every* $\epsilon > 0$.

The corollary follows easily from the previous results using the bias-variance decomposition. Details are given in Section 7. In case $d = O(1)$, then the result can be strengthened to

$$\mathcal{R}(h^\star) = \widetilde{O}_P\left(n^{-4/(4+r)}\right) \tag{5.18}$$

so that the optimal rate of convergence is obtained up to logarithmic factors, as if the relevant variables were known in advance and isolated.

Assumption (A1) imposes a condition on the sampling density of the irrelevant variables, requiring that it is sufficiently smooth in the irrelevant dimensions. Using Theorem 5.1, we can modify the threshold used by the algorithm in order to remove this assumption.

**Theorem 5.5.** *Consider a modified rodeo algorithm where the threshold is changed to*

$$\lambda_j = \rho_n \beta^{3t} + s_j \sqrt{2 \log(n c_n)} \tag{5.19}$$

*at step* $t$, *where* $\rho_n > 0$ *satisfies*

$$\frac{\rho_n}{h_0^3} \longrightarrow \infty, \quad \frac{\rho_n}{h_0} \longrightarrow 0 \quad \text{as } n \longrightarrow \infty. \tag{5.20}$$

*Then the results of Theorem 5.3 hold under assumption (A2), without assuming (A1).*

# VI. Extensions and Variations of the Rodeo

The rodeo represents a general strategy for nonparametric estimation, based on the idea of regularizing or testing the derivatives of an estimator with respect to smoothing parameters. There are many ways in which this basic strategy can be realized. In this section we discuss several variants of the basic hard thresholding version of the rodeo, including a soft thresholding version, a global rather than local bandwidth selection procedure, the use of testing and generalized cross validation, and connections to least angle regression. Further numerical examples are also given to illustrate these ideas.

### *A. Subtracting off a Linear Lasso*

Local linear regression is a nonparametric method that contains linear regression as a special case when $h \to \infty$. If the true function is linear but only a subset of the variables are relevant, then the rodeo will fail to separate the relevant and irrelevant variables since relevance is defined in terms of departures from the limiting parametric model. Indeed, the results depend on the Hessian of $m$ which is zero in the linear case. The rodeo may return a full linear fit with all variables. A simple modification fixes this problem. First, do linear variable selection using, say, the lasso (Tibshirani, 1996). Then run the rodeo on the residuals from that fit, but using all of the variables. An example of this procedure is given below in Section 6.D.

### *B. Other Estimators and Other Paths*

We have taken the estimate

$$\widehat{D}_j(h) = Z_j(h) I(|Z_j(h)| > \lambda_j) \tag{6.2}$$

with the result that

$$\widetilde{m}(x) = \widehat{m}_{h_0}(x) - \int_0^1 \langle \widehat{D}(s), \dot{h}(s) \rangle ds = \widehat{m}_{h\star}(x). \tag{6.3}$$

There are many possible generalizations. First, we can replace $\widehat{D}$ with the soft-thresholded estimate

$$\widehat{D}_j(t) \;=\; \text{sign}(Z_j(h))\,(|Z_j(h)| - \lambda_j)_+ \tag{6.4}$$

where the index $t$ denotes the $t^{\text{th}}$ step of the algorithm. Since $h_j$ is updated multiplicatively as $h_j \leftarrow \beta h_j$, the differential $dh_j(t)$ is given by $dh_j(t) = (1-\beta)h_j$. Using the resulting estimate of $D(t)$ and finite difference approximation for $\dot{h}(t)$ leads to the algorithm detailed in Figure 7.

*Rodeo: Soft thresholding version*

1. *Select* parameter $0 < \beta < 1$ and initial bandwidth $h_0$, satisfying $1 \leq h_0 \leq \log^{\ell/d} n$, for a fixed constant $\ell$. Let $c_n$ be a sequence satisfying $c_n = O(1)$.

2. *Initialize* the bandwidths, and activate all covariates:

   (a) $h_j = h_0$, $j = 1, 2, \ldots, d$.

   (b) $\mathcal{A} = \{1, 2, \ldots, d\}$

   (c) Initialize step, $t = 1$.

3. *While* $\mathcal{A}$ *is nonempty*

   (a) Set $dh_j(t) = 0$, $j = 1, \ldots, d$.

   (b) Do for each $j \in \mathcal{A}$:

      (1) Compute the estimated derivative expectation $Z_j$ and $s_j$.

      (2) Compute the threshold $\lambda_j = s_j \sqrt{2 \log(nc_n)}$.

      (3) If $|Z_j| > \lambda_j$, set $dh_j(t) = (1 - \beta)\, h_j$ and $h_j \leftarrow \beta h_j$; otherwise remove $j$ from $\mathcal{A}$.

      (4) Set $\widehat{D}_j(t) = \text{sign}(Z_j(h))\, (|Z_j(h)| - \lambda_j)_+$.

   (c) Increment step, $t \leftarrow t + 1$.

4. *Output* bandwidths $h^\star = (h_1, \ldots, h_d)$ and estimator

$$\widetilde{m}(x) = \widehat{m}_{h_0}(x) - \sum_{s=1}^{t} \big\langle \widehat{D}(s), dh(s) \big\rangle \tag{6.1}$$

Figure 7: The soft thresholding version of the rodeo.

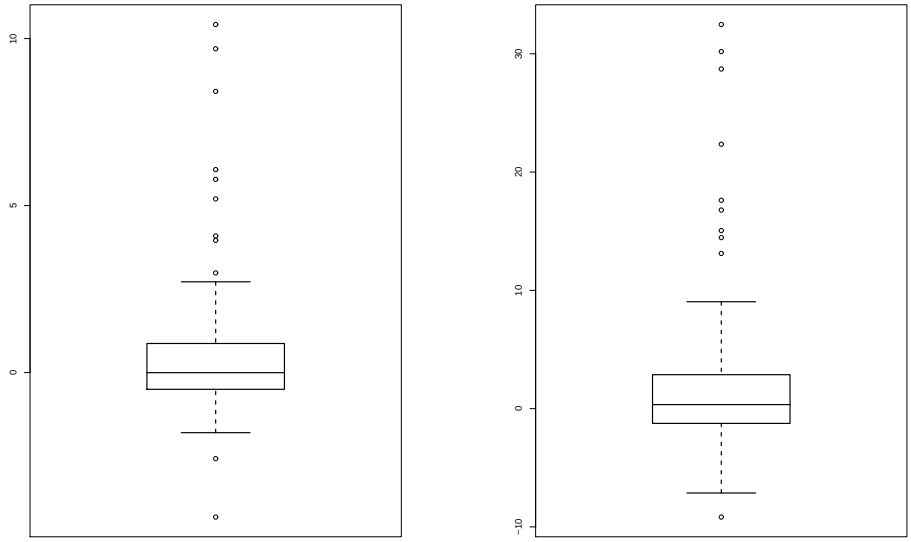


Figure 8: Comparison of hard and soft thresholding. Left: $m(x) = 5x_1^2x_2^2$, $d = 10$ and $\sigma = 0.5$; right: $m(x) = 2(x_1 + 1)^3 + 2\sin(10\,x_2)$, $d = 10$ and $\sigma = 1$. The hard and soft thresholding versions of the rodeo were compared on 100 randomly generated datasets, with a single random test point $x$ chosen for each; $\beta = 0.9$. The plots show two views of the difference of losses, $(\widetilde{m}_{\text{hard}}(x) - m(x))^2 - (\widetilde{m}_{\text{soft}}(x) - m(x))^2$; positive values indicate an advantage for soft thresholding.

Figure 8 shows a comparison of the hard and soft thresholding versions of the rodeo on the example function $m(x) = 2(x_1+1)^3+2\sin(10\,x_2)$ in $d = 10$ dimensions with $\sigma = 1$; $\beta$ was set to 0.9. For each of 100 randomly generated datasets, a random test point $x \sim \text{Uniform}(0,1)^d$ was generated, and the difference in losses was computed:

$$(\widetilde{m}_{\text{hard}}(x) - m(x))^2 - (\widetilde{m}_{\text{soft}}(x) - m(x))^2 \ . \tag{6.5}$$

Thus, positive values indicate an advantage for soft thresholding, which is seen to be slightly more robust on this example.

Another natural extension would be to consider more general paths than paths that are restricted to be parallel to the axes. We leave this direction to future work.

*C. Global Rodeo*

We have focused on estimation of $m$ locally at a point $x$. The idea can be extended to carry out global bandwidth and variable selection by averaging over multiple evaluation points $x_1, \ldots, x_k$. These could be points interest for estimation, could be randomly chosen, or could be taken to be identical to the observed $X_i$s.

Averaging the $Z_j$s directly leads to a statistic whose mean for relevant variables is asymptotically $k^{-1}h_j \sum_{i=1}^{k} m_{jj}(x_i)$. Because of sign changes in $m_{jj}(x)$, cancellations can occur resulting in a small value of the statistic for relevant variables. To eliminate the sign cancellation, we square the statistic. Another way of deriving a global method would be to use the statistic $\sup_x |Z_j^*(x)|$.

Let $x_1, \ldots, x_k$ denote the evaluation points. Let

$$Z_j(x_i) = \sum_{s=1}^{n} Y_s G_j(X_s, x_i). \tag{6.6}$$

Then define the statistic

$$T_j \equiv \frac{1}{k} \sum_{i=1}^{k} Z_j^2(x_i) = \frac{1}{k} Y^T P_j Y \tag{6.7}$$

where $P_j = \mathcal{G}_j \mathcal{G}_j^T$, with $\mathcal{G}_j(s,i) = G_j(X_s, x_i)$.

If $j \in R^c$ then we have $\mathbb{E}(Z_j(x_i)) = o(1)$, so it follows that, conditionally,

$$\mathbb{E}(T_j) = \frac{\sigma^2}{k} \operatorname{tr}(P_j) + o_P(1) \tag{6.8a}$$

$$\mathbb{V}(T_j) = \frac{2\sigma^4}{k^2} \operatorname{tr}(P_j P_j) + o_P(1). \tag{6.8b}$$

We take the threshold to be

$$\lambda_j = \frac{\widehat{\sigma}^2}{k} \operatorname{tr}(P_j) + 2\frac{\widehat{\sigma}^2}{k} \sqrt{\operatorname{tr}(P_j P_j) \log(c_n n)}\ . \tag{6.9}$$

Note that if $j > r$, we have

$$\mathbb{E}(T_j) = \frac{1}{k} \sum_i s_j^2(X_i) + O(h_0^6) \tag{6.10}$$

but for $j \leq r$ we have

$$\mathbb{E}(T_j) = \frac{1}{k} \sum_i s_j^2(X_i) + O(h_0^2). \tag{6.11}$$

We give an example of this algorithm in the following section, leaving the detailed analysis of the asymptotics of this estimator to future work.

*D. Greedier Rodeo and LARS*

The rodeo is related to least angle regression (LARS) (Efron et al., 2004). In forward stagewise linear regression, one performs variable selection incrementally. LARS gives a refinement where at each step in the algorithm, one adds the covariate most correlated with the residuals of the current fit, in small, incremental steps. LARS takes steps of a particular size: the smallest step that makes the largest correlation equal to the next-largest correlation. Efron et al. (2004) show that the lasso can be obtained by a simple modification of LARS.

The rodeo can be seen as a nonparametric version of forward stagewise regression. Note first that $Z_j$ is essentially the correlation between the $Y_i$s and the $G_j(X_i, x, h)$s (the change in the effective kernel). Reducing the bandwidth is like adding in more of that variable. Suppose now that we make the following modifications to the rodeo: (i) change the bandwidths one at a time, based on the largest $Z_j^* = Z_j/\lambda_j$, (ii) reduce the bandwidth continuously, rather than in discrete steps, until the largest $Z_j$ is equal to the next largest. This version can then be thought of as a nonparametric formulation of LARS.

In fact, we can go further and embed the rodeo within LARS to get a fast nonparametric method. We do this by replacing the derivatives of the fit in the rodeo with differences. Then we iterate variable selection with bandwidth selection.

- Set $h = (h_0, \ldots, h_0)$. Define $d$-dimensional pseudo-covariates $\widetilde{X}_i$, $i = 1, \ldots, n$, by

$$\widetilde{X}_i(j) = G_j(X_i, x, h), \quad j = 1, \ldots, d. \tag{6.12}$$

  Now run the LARS algorithm, regressing the $Y_i$'s on the pseudo-covariates, up to some pre-defined stopping point. This step essentially chooses relevant variables at the resolution of the starting bandwidth $h = (h_0, \ldots, h_0)$.

- Define new pseudo-covariates $\widetilde{X}_i$, $i = 1, \ldots, n$, by

$$\widetilde{X}_i(j) = G_j(X_i, x, h') - G_j(X_i, x, h), \quad j = 1, \ldots, d \tag{6.13}$$

  where $h'$ has $h_j$ replaced by $\beta h_j$. Note that adding the $j^{\text{th}}$ covariate corresponds to reducing the bandwidth from $h_j$ to $\beta h_j$. Now run the LARS algorithm up to some pre-defined stopping point.

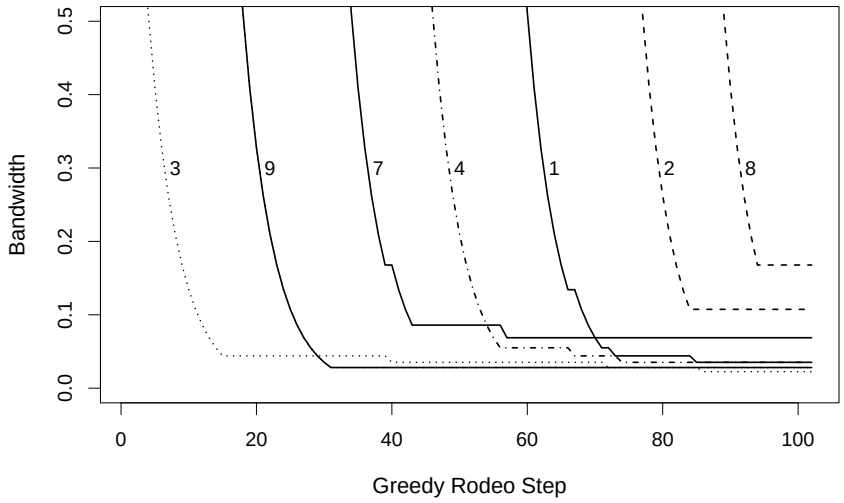


Figure 9: Greedy rodeo on the diabetes data, used to illustrate LARS (Efron et al., 2004). A set of $k = 100$ of the total $n = 442$ points were sampled ($d = 10$), and the bandwidth for the variable with largest average $|Z_j|/\lambda_j$ was reduced in each step.

- Repeat the last step until a stopping criterion is satisfied.

One advantage of this method is that it can be implemented using existing LARS software. We leave the development of the theory for this approach to future work. Some examples of the greedy version of this algorithm follow.

*D.1 Diabetes example*

Figure 9 shows the result of running the greedy version of the rodeo on the diabetes dataset used by Efron et al. (2004) to illustrate LARS . The algorithm averages $Z_j^*$ over a randomly chosen set of $k = 100$ data points, and reduces the bandwidth for the variable with the largest value; note that no estimate of $\sigma$ is required. The resulting variable ordering is seen to be very similar to, but different from, the ordering obtained from the parametric LARS fit. The variables were selected in the order 3 (body mass index), 9 (serum), 7 (serum), 4 (blood pressure), 1 (age), 2 (sex), 8 (serum), 5 (serum), 10 (serum), 6 (serum). The LARS algorithm adds variables in the order 3, 9, 4, 7, 2, 10, 5, 8, 6, 1. One notable difference is in the position of the age variable.

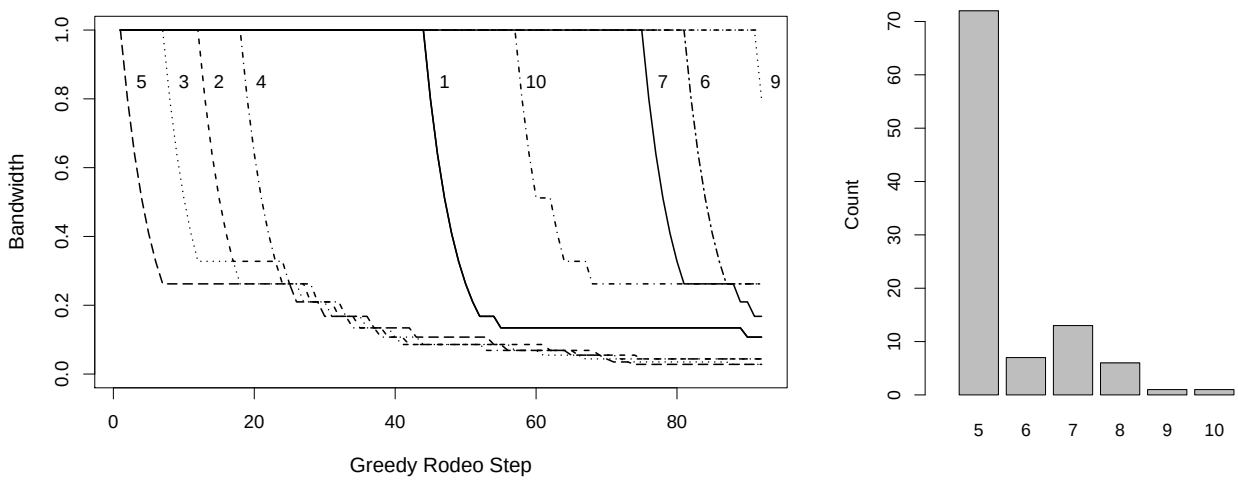


Figure 10: Left: A typical run of the greedy algorithm on Turlach's example. The bandwidths are first reduced for variables $x_2, x_3, x_4, x_5$, and then the relevant, but uncorrelated with $Y$ variable $x_1$ is added to the model; the irrelevant variables enter the model last. Right: Histogram of the position at which variable $x_1$ is selected, over 100 runs of the algorithm

*D.2 Turlach's example*

In the discussion to the LARS paper, Berwin Turlach (Turlach, 2004) gives an interesting example of where LARS and the lasso fails. The function is

$$Y = \left(X_1 - \frac{1}{2}\right)^2 + X_2 + X_3 + X_4 + X_5 + \varepsilon \tag{6.14}$$

with ten variables $X_i \sim \text{Uniform}\,(0, 1)$ and $\sigma = 0.05$. Although $X_1$ is a relevant variable, it is uncorrelated with $Y$, and LARS and the lasso miss it.

Figure 10 shows the greedy algorithm on this example, where bandwidth corresponding to the largest average $Z_j^*$ is reduced in each step. We use kernel regression rather than local linear regression as the underlying estimator, without first subtracting off a Lasso fit. The variables $x_2$, $x_3$, $x_4$, $x_5$ are linear in the model, but are selected first in every run. Variable $x_1$ is selected fifth in 72 of the 100 runs; a typical run of the algorithm is shown in the left plot. In contrast, as discussed in Turlach (2004), LARS selects $x_1$ in position 5 about 25% of the time.

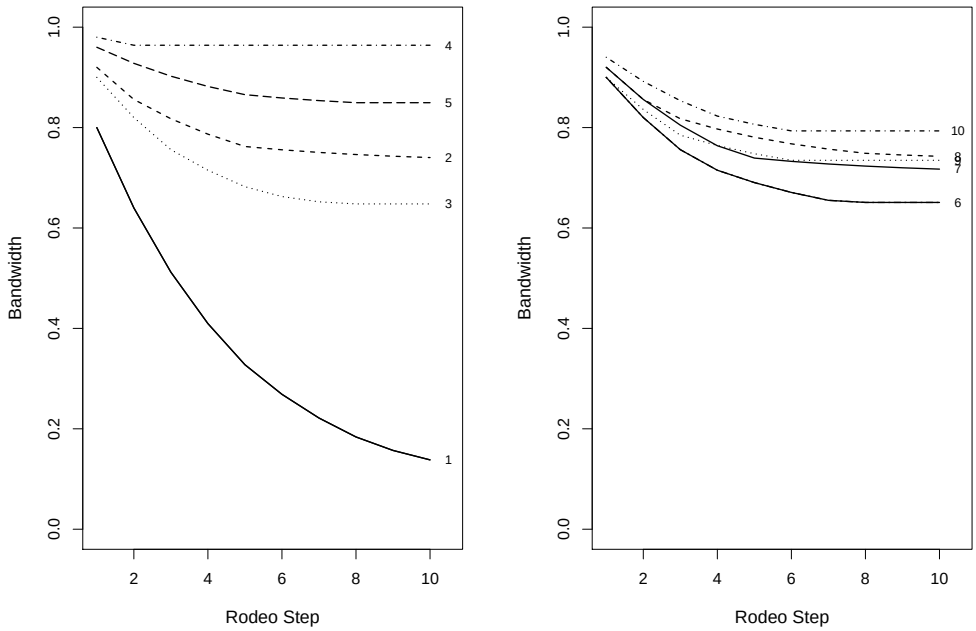


Figure 11: The global rodeo averaged over 10 runs on Turlach's example. The left plot shows the bandwidths for the five relevant variables. Since the linear effects (variables two through five) have been subtracted off, bandwidths $h_2, h_3, h_4, h_5$ are not shrunk. The right plot shows the bandwidths for the other, irrelevant, variables.

Figure 11 shows bandwidth traces for this example using the global algorithm described in Section 6.C with $k = 20$ evaluation points randomly subselected from the data, and $\sigma$ taken to be known. Before starting the rodeo, we subtract off a linear least squares fit, and run the rodeo on the residuals. The first plot shows $h_1, \ldots, h_5$. The lowest line is $h_1$ which shrinks the most since $m$ is a nonlinear function of $x_1$. The other curves are the linear effects. The right plot shows the traces for $h_6, \ldots, h_{10}$, the bandwidths for the irrelevant variables.

## VII. Proofs of Technical Results

In this section we give the proofs of the results stated in Section 5.

*A. Proof of Theorem 5.1*

We follow the setup of Ruppert and Wand (1994) except for two differences: the irrelevant variables have different leading terms in the expansions than relevant variables and $d$ is increasing; this requires a more detailed asymptotic analysis than is carried out by Ruppert and Wand (1994).

Let $\mathcal{H} = \begin{pmatrix} \mathcal{H}_R & 0 \\ 0 & 0 \end{pmatrix}$ denote the Hessian of $m(x)$, let $h_j^{(t)}$ denote the $j^{\text{th}}$ bandwidth at step $t$ and denote the bandwidth matrix by $H^{(t)} = \text{diag}((h_1^{(t)})^2, \ldots, (h_d^{(t)})^2)$. Similarly, let $H_R^{(t)} = \text{diag}((h_1^{(t)})^2, \ldots, (h_r^{(t)})^2)$. Define

$$\mu_j^{(t)} = \frac{\partial}{\partial h_j} \mathbb{E}\left[\widehat{m}_{H^{(t)}}(x) - m(x) \,|\, X_1, \ldots, X_d\right] \tag{7.1}$$

In the following we sometimes supress the superscript $t$.

Let $\nabla m$ be the gradient of $m$ at $x$, and let

$$Q = ((X_1 - x)^T \mathcal{H}(X_1 - x), \ldots, (X_n - x)^T \mathcal{H}(X_n - x))^T. \tag{7.2}$$

Note that $\nabla m$ and $Q$ are only functions of the relevant variables. Then

$$m(X_i) = m(x) + (X_i - x)^T \nabla m + \frac{1}{2} Q_i + R_i \tag{7.3}$$

where, using multi-index notation,

$$R_i = \frac{1}{6} \sum_{|\alpha|=3} (X_i - x)^\alpha \int_0^1 D^\alpha m\left((1-s)x + sX_i\right) \, ds = \frac{1}{6} \sum_{|\alpha|=3} (X_i - x)^\alpha D^\alpha m(\xi_{i,\alpha}) \tag{7.4}$$

where $\xi_{i,\alpha}$ is some point between $X_i$ and $x$.

So, with $M = (m(X_1), \ldots, m(X_n))^T$,

$$M = X_x \begin{pmatrix} m(x) \\ \nabla m \end{pmatrix} + \frac{1}{2} Q + R \tag{7.5}$$

where $R = (R_1, \ldots, R_n)^T$. Since $S_x X_x (m(x), \nabla m)^T = m(x)$, the bias $b(x) = \mathbb{E}(\widehat{m}_H(x)) - m(x)$ is given by

$$\begin{aligned} b(x) &= e_1^T S_x M - m(x) = \frac{1}{2} S_x Q + S_x R & \text{(7.6a)} \\ &= \frac{1}{2} e_1^T (X_x^T W_x X_x)^{-1} X_x^T W_x Q + e_1^T (X_x^T W_x X_x)^{-1} X_x^T W_x R & \text{(7.6b)} \\ &= \frac{1}{2} e_1^T \Upsilon_n^{-1} \Gamma_n + e_1^T \Upsilon_n^{-1} \frac{1}{n} X_x^T W_x R & \text{(7.6c)} \end{aligned}$$

where $\Upsilon_n = n^{-1}(X_x^T W_x X_x)$ and $\Gamma_n = n^{-1}(X_x^T W_x Q)$.

*Analysis of* $\Upsilon_n$. We write

$$\Upsilon_n \;=\; \begin{pmatrix} \frac{1}{n}\sum_{i=1}^n W_i & \frac{1}{n}\sum_{i=1}^n W_i(X_i - x)^T \\ \frac{1}{n}\sum_{i=1}^n W_i(X_i - x) & \frac{1}{n}\sum_{i=1}^n W_i(X_i - x)(X_i - x)^T \end{pmatrix} = \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix} \tag{7.7a}$$

where

$$W_i \;=\; \prod_{j=1}^d \frac{1}{h_j} K\left(\frac{x_j - X_{ij}}{h_j}\right) \;=\; \prod_{j=1}^d K_{h_j}(x_j - X_{ij}) \tag{7.8}$$

Now, the variance of $W_i$ can be bounded as

$$\mathbb{V}\mathrm{ar}(W_i) \le E(W_i^2) = \prod_{j=1}^d \frac{1}{h_j} \mathbb{E}\left[\prod_{j=1}^d \frac{1}{h_j} K^2\left(\frac{x_j - X_{ij}}{h_j}\right)\right] \le C_K^d \prod_{j=1}^d \frac{1}{h_j} \tag{7.9}$$

where $C_K = \sup_u K(u)$. Therefore,

$$\mathbb{V}\mathrm{ar}(A_{11}) \;\le\; \frac{C_K^d}{n} \prod_{j=1}^d \frac{1}{h_j} \le \frac{1}{n} \frac{C_K^d}{\beta^{ds} h_0^d} \equiv \frac{\Delta_s}{n}. \tag{7.10}$$

Also, $W_i \le \Delta_s$. Note that $\beta^{T_n d} n h_0^d / C_k^d = \Omega\left(n^{1-\alpha-\log\log\log n/\log\log n}\right)$, where $\alpha$ is the constant appearing in (5.4). Hence, by Bernstein's inequality, for any $\delta > 0$,

$$\begin{aligned} \mathbb{P}(|A_{11} - \mathbb{E}(A_{11})| > \epsilon) \;&\le\; 2\exp\left\{-\frac{1}{2}\left(\frac{n\epsilon^2}{\Delta + \Delta\epsilon/3}\right)\right\} && \text{(7.11a)} \\ &=\; 2\exp\left\{-\frac{1}{2}\left(\frac{n\epsilon^2 \beta^{dt} h_0^d}{C_K^d(1+\epsilon/3)}\right)\right\} && \text{(7.11b)} \\ &\le\; 2\exp\left\{-\frac{1}{2}\left(\frac{n\epsilon^2 \beta^{dT_n} h_0^d}{C_K^d(1+\epsilon/3)}\right)\right\} && \text{(7.11c)} \\ &\le\; 2\exp\left\{-\frac{1}{2} n^{(1-\delta)}\epsilon^2\right\}. && \text{(7.11d)} \end{aligned}$$

Therefore, uniformly for all $t$ and $j$, $A_{11}^{(t)} = \mathbb{E}(A_{11}^{(t)}) + O_P(1/n^{(1/2)-\delta})$ for any $\delta > 0$. Let $D_2$

denote the Hessian of $f$. For some $\widetilde{v}$ between $x$ and $v$,

$$\begin{aligned}
\mathbb{E}(A_{11}) &= \int \frac{1}{h_1 h_2 \cdots h_d} K(H^{-1/2}(x-u)) f(u)\, du && (7.12a)\\
&= \int K(v) f(x - H^{1/2}v)\, dv && (7.12b)\\
&= f(x) - \int K(v)(H^{1/2}v)^T \nabla f(x)\, dv + \frac{1}{2}\int K(v)(H^{1/2}v)^T D_2(\widetilde{v}) H^{1/2} v\, dv && (7.12c)\\
&= f(x) + \frac{1}{2}\int K(v)(H^{1/2}v)^T D_2(\widetilde{v}) H^{1/2} v\, dv && (7.12d)\\
&= f(x) + r_1 && (7.12e)
\end{aligned}$$

where from (5.6)

$$r_1 = \frac{1}{2}\int K(v)(H^{1/2}v)^T D_2(\widetilde{v}) H^{1/2} v\, dv = O\left(d\ h_0^2\ \max_{a,b}\sup_x \left|\frac{\partial^2 f(x)}{\partial x_a \partial x_b}\right|\right) = O\left((h_j^{(t)})^2/d\right) \tag{7.13}$$

Therefore we have, uniformly over $1 \le t \le T_n$,

$$A_{11} = f(x) + O\left((h_j^{(t)})^2/d\right). \tag{7.14}$$

Next consider $A_{21}$. By a similar argument as above, uniformly for all $t$ and $j$, $A_{21}^{(t)} = \mathbb{E}(A_{21}^{(t)}) + O_P(1/n^{(1/2)-\delta})$ for any $\delta > 0$. Also,

$$\begin{aligned}
\mathbb{E}(A_{21}) &= \int \frac{1}{h_1 h_2 \cdots h_d} K(H^{-1/2}(x-u))(u-x) f(u)\, du && (7.15a)\\
&= \int K(v) H^{1/2} v f(x + H^{1/2}v)\, dv && (7.15b)\\
&= f(x)\underbrace{\int K(v) H^{1/2} v\, dv}_{0} + \int K(v) H^{1/2} v (H^{1/2}v)^T D\, dv\\
&\quad + \frac{1}{2}\underbrace{\int K(v) H^{1/2} v (H^{1/2}v)^T D_2 H^{1/2} v\, dv}_{0}\\
&\quad + \frac{1}{6}\int K(v) H^{1/2} v \sum_{abc} h_a h_b h_c v_a v_b v_c D_3(a,b,c)\, dv && (7.15c)
\end{aligned}$$

where $D_2$ and $D_3$ are the arrays of second and third derivatives of $f$, the latter being evaluated

at some point between $x$ and $u$. So,

$$\mathbb{E}(A_{21}) = \int K(v)H^{1/2}v(H^{1/2}v)^T D\,dv + \frac{1}{6}\int K(v)H^{1/2}v\sum_{abc} h_a h_b h_c v_a v_b v_c D_3(a,b,c)\,dv \tag{7.16a}$$

$$= \nu_2(K)HD + r_2 \tag{7.16b}$$

where

$$r_2 = \frac{1}{6}\int K(v)H^{1/2}v\sum_{abc} h_a h_b h_c v_a v_b v_c D_3(a,b,c)\,dv \tag{7.17a}$$

$$= O\left(h_0^4\, d^2 \max_{a,b,c,d}\sup_x \left|\frac{\partial^3 f(x)}{\partial x_a \partial x_b \partial x_c}\right| \mathbf{1}\right) = O\left(((h_j^{(t)})^4/d)\,\mathbf{1}\right) \tag{7.17b}$$

using (5.7) and the definition of $\nu_2$ in equation (3.5). Therefore, $A_{21} = \nu_2(K)HD + O_P\left(((h_j^{(t)})^4/d)\mathbf{1}\right)$.

Now we turn to $A_{22}$. Again we have convergence to its mean and

$$\mathbb{E}(A_{22}) = \int K(v)H^{1/2}v(H^{1/2}v)^T f(x + H^{1/2}v)\,dv \tag{7.18a}$$

$$= f(x)\int K(v)H^{1/2}v(H^{1/2}v)^T\,dv + \frac{1}{2}\int K(v)H^{1/2}v(H^{1/2}v)^T(H^{1/2}v)^T D\,dv$$
$$+\frac{1}{6}\int K(v)H^{1/2}v(H^{1/2}v)^T(H^{1/2}v)^T D_2(H^{1/2}v) \tag{7.18b}$$

$$= \nu_2 f(x)H + R_3 \tag{7.18c}$$

where $R_3$ is a matrix whose off-diagonal elements are of order

$$O\left(h_0^4 \max_{a,b}\sup_x \left|\frac{\partial^2 f(x)}{\partial x_a \partial x_b}\right|\right) = O\left((h_j^{(t)})^4/d^2\right) \tag{7.19}$$

and whose diagonal elements are of order

$$O\left(d h_0^4 \max_{a,b}\sup_x \left|\frac{\partial^2 f(x)}{\partial x_a \partial x_b}\right|\right) = O\left((h_j^{(t)})^4/d\right). \tag{7.20}$$

Thus,

$$\Upsilon_n = \begin{pmatrix} f(x) + r_1 & \nu_2 D^\top H + r_2^T \\ \nu_2 H^T D + r_2 & \nu_2 f(x)H + R_3 \end{pmatrix}. \tag{7.21}$$

We can now write

$$\Upsilon_n = A + vv^T + E \tag{7.22}$$

where

$$A = \begin{pmatrix} f(x) - 1 & 0 \\ 0 & \nu_2 f(x) H \end{pmatrix}, \tag{7.23}$$

$$v^\top = (1,\ \nu_2 D^\top H) \tag{7.24}$$

and

$$E = \begin{pmatrix} r_1 & r_2^T \\ r_2 & R_3 - \nu_2^2 H D D^T H \end{pmatrix}. \tag{7.25}$$

To invert the matrix we use the following result. Let $B$ be an invertible matrix and let $\Delta$ be any matrix such that $|\Delta| \le \epsilon |B|$ (interpreted elementwise). Then,

$$|(B+\Delta)^{-1} - B^{-1}| \le \epsilon |B^{-1}|\ |B|\ |B^{-1}| + O(\epsilon^2) \tag{7.26}$$

where the absolute values and inequalities are elementwise. Thus,

$$|\Upsilon_n^{-1} - (A + vv^T)^{-1}| \le \epsilon |(A+vv^T)^{-1}|\ |A+vv^T|\ |(A+vv^T)^{-1}| + O(\epsilon^2) \tag{7.27}$$

where $\epsilon = \max_{ij} |E_{ij}|$. It now follows from (7.13), (7.17b), (7.19), (7.20) and a union bound that

$$\max_{j,k} |\Upsilon_n^{-1}(j,k) - (A+vv^T)^{-1}(j,k)| = o_P(h_0^2). \tag{7.28}$$

To compute $(A + vv^T)^{-1}$ we use the matrix inversion lemma (Woodbury formula)

$$(A + vv^\top)^{-1} = A^{-1} - A^{-1} v (1 + v^\top A^{-1} v)^{-1} v^\top A^{-1}\,. \tag{7.29}$$

Now,

$$A^{-1} = \begin{pmatrix} \frac{1}{f(x)-1} & 0 \\ 0 & \frac{H^{-1}}{\nu_2 f(x)} \end{pmatrix}, \tag{7.30}$$

and

$$1 + v^\top A^{-1} v \ =\ \frac{f(x)}{f(x)-1} + \frac{\nu_2 D^\top H D}{f(x)} = \frac{f(x)}{f(x)-1}\left(1 + O(dh_0^2 \max D_2)\right) \tag{7.31}$$

so that

$$(1 + v^\top A^{-1} v)^{-1} \ =\ \frac{f(x)-1}{f(x)}(1 + o(1)). \tag{7.32}$$

Also,

$$A^{-1} v v^\top A^{-1} \ =\ \begin{pmatrix} \frac{1}{(f(x)-1)^2} & \frac{D^\top}{f(x)(f(x)-1)} \\ \frac{D}{f(x)(f(x)-1)} & \frac{DD^T}{f(x)^2} \end{pmatrix}. \tag{7.33}$$

Finally, we have that

$$\Upsilon_n^{-1} = (A + vv^\top)^{-1} + o_P(1) = \begin{pmatrix} \frac{1}{f(x)} & -\frac{D^T}{f^2(x)} \\ -\frac{D}{f^2(x)} & \frac{H^{-1}}{\nu_2 f(x)} - \frac{DD^T(f(x)-1)}{f^3(x)} \end{pmatrix} + o_P(h_0^2). \tag{7.34}$$

This completes the analysis of $\Upsilon_n$.

*Analysis of* $\Gamma_n = \frac{1}{n} X_x^T W_x Q$. We can write

$$\Gamma_n \equiv \frac{1}{n} X_x^T W_x Q = \begin{pmatrix} \frac{1}{n}\sum_{i=1}^n W_i (X_i - x)^T \mathcal{H}(X_i - x) \\ \frac{1}{n}\sum_{i=1}^n (X_i - x) W_i (X_i - x)^T \mathcal{H}(X_i - x) \end{pmatrix} = \begin{pmatrix} \gamma_1 \\ \gamma_2 \end{pmatrix}. \tag{7.35}$$

Now,

$$\begin{aligned}
\mathbb{E}(\gamma_1) &= \int K(v)(H^{1/2}v)^T \mathcal{H}(H^{1/2}v) f(x + H^{1/2}v)\, dv && (7.36a)\\
&= f(x) \int K(v)(H^{1/2}v)^T \mathcal{H}(H^{1/2}v)\, dv + \int K(v)(H^{1/2}v)^T \mathcal{H}(H^{1/2}v) D^T (H^{1/2}v)\, dv \\
&\quad + \frac{1}{2} \int K(v)(H^{1/2}v)^T \mathcal{H}(H^{1/2}v)(H^{1/2}v)^T D_2^T (H^{1/2}v)\, dv && (7.36b)\\
&= \nu_2 f(x) \operatorname{tr}(H\mathcal{H}_R) + O(drh_0^4 \max_{ab} D_2(a,b)) && (7.36c)\\
&= \nu_2 f(x) \operatorname{tr}(H\mathcal{H}_R) + O((h_j^{(t)})^4/d). && (7.36d)
\end{aligned}$$

The stochastic analysis of $\gamma_1$ from its mean is similar to the analysis of $A_{12}$ and we have $\gamma_1 = \nu_2 f(x) \operatorname{tr}(H\mathcal{H}_R) + o_P\left((h_j^{(t)})^4/d\right)$.

Next,

$$\begin{aligned}
\mathbb{E}(\gamma_2) &= \int K(v)(H^{1/2}v)(H^{1/2}v)^T \mathcal{H}(H^{1/2}v) f(x + H^{1/2}v)\, dv && (7.37a)\\
&= f(x) \int (H^{1/2}v) K(v)(H^{1/2}v)^T \mathcal{H}(H^{1/2}v)\, dv \\
&\quad + \int K(v)(H^{1/2}v)(H^{1/2}v)^T \mathcal{H}(H^{1/2}v) D^T (H^{1/2}v)\, dv \\
&\quad + \frac{1}{2} \int K(v)(H^{1/2}v)(H^{1/2}v)^T \mathcal{H}(H^{1/2}v)(H^{1/2}v)^T D_2^T (H^{1/2}v)\, dv && (7.37b)\\
&= \int K(v)(H^{1/2}v)(H^{1/2}v)^T \mathcal{H}(H^{1/2}v) D^T (H^{1/2}v)\, dv \\
&\quad + \frac{1}{2} \int K(v)(H^{1/2}v)(H^{1/2}v)^T \mathcal{H}(H^{1/2}v)(H^{1/2}v)^T D_2^T (H^{1/2}v)\, dv && (7.37c)\\
&= \Lambda + O(drh_0^5 \max_{ab} D_2(a,b)\, \mathbf{1}) && (7.37d)\\
&= \Lambda + O\left((h_0^5/d)\mathbf{1}\right) && (7.37e)
\end{aligned}$$

where $\Lambda$ is a $d$-vector with $j^{\text{th}}$ component

$$\Lambda_j = h_j^4 \nu_4 \mathcal{H}_{jj} f_j + I(j \le r) \nu_2^2 h_j^2 \sum_{k=1}^{r} h_k^2 \mathcal{H}_{jk} f_k + \nu_2^2 h_j^2 f_j \sum_{k=1}^{r} h_k^2 \mathcal{H}_{kk}. \tag{7.38}$$

Thus,

$$\Gamma_n = \begin{pmatrix} \nu_2 f(x) \operatorname{tr}(H \mathcal{H}_R) \\ \Lambda \end{pmatrix} + O_P\left((h_0^5/d)\mathbf{1}\right). \tag{7.39}$$

*Analysis of remainder* $e_1^T \Upsilon_n^{-1} \frac{1}{n} X_x^T W_x R$. We can write

$$\frac{1}{n} X_x^T W_x R \;=\; \begin{pmatrix} \frac{1}{n}\sum_{i=1}^n W_i R_i \\ \frac{1}{n}\sum_{i=1}^n (X_i - x) W_i R_i \end{pmatrix} = \begin{pmatrix} \delta_1 \\ \delta_2 \end{pmatrix}. \tag{7.40}$$

Let $m_3$ denote the array of third derivatives of $m$ and let $\widetilde{v}_j = h_j v_j$. Then,

$$\begin{aligned}
\mathbb{E}(\delta_1) &= \int K(v)\, m_3(\xi) \sum_{1 \le a,b,c \le r} \widetilde{v}_a \widetilde{v}_b \widetilde{v}_c f(x + \widetilde{v})\, dv && (7.41a)\\
&= f(x) \int K(v)\, m_3(\xi) \sum_{1 \le a,b,c \le r} \widetilde{v}_a \widetilde{v}_b \widetilde{v}_c\, dv + \int K(v)\, m_3(\xi) \sum_{1 \le a,b,c \le r} \widetilde{v}_a \widetilde{v}_b \widetilde{v}_c\, (\widetilde{v}^T D)\, dv && (7.41b)\\
&= m_3(x) \int K(v) \sum_{1 \le a,b,c \le r} \widetilde{v}_a \widetilde{v}_b \widetilde{v}_c\, (\widetilde{v}^T D)\, dv + o\left((h_j^{(t)})^4\right) && (7.41c)
\end{aligned}$$

and so

$$|\mathbb{E}(\delta_1)| = O\left(\sup_x |m_3(x)|\; \sup_x |D(x)| \sum_{j,k \le r} h_j^2 h_k^2\right) = O\left(\sum_{j,k \le r} h_j^2 h_k^2\right) \tag{7.42}$$

Similarly,

$$\begin{aligned}
\mathbb{E}(\delta_2) &= \int K(v)\, m_3(\xi)\, \widetilde{v} \sum_{1 \le a,b,c \le r} \widetilde{v}_a \widetilde{v}_b \widetilde{v}_c f(x + \widetilde{v})\, dv && (7.43a)\\
&= f(x) \int K(v)\, m_3(\xi)\, \widetilde{v} \sum_{1 \le a,b,c \le r} \widetilde{v}_a \widetilde{v}_b \widetilde{v}_c\, dv + o\left((h_j^{(t)})^4\right) && (7.43b)\\
&= O\left(\sup_x |m_3(x)| \sum_{j,k \le r} h_j^2 h_k^2\right). && (7.43c)
\end{aligned}$$

Hence,

$$e_1^T \Upsilon_n^{-1} \frac{1}{n} X_x^T W_x R = O_P\left(\sum_{j,k \le r} h_j^2 h_k^2\right). \tag{7.44}$$

Putting all of this together, we conclude that the leading term of $b(x)$ is

$$\frac{1}{2}\nu_2 \operatorname{tr}(H\mathcal{H}_R) - \frac{1}{2}\frac{D^T\Lambda}{f^2(x)} \tag{7.45}$$

where $\Lambda$ is given in (7.38). Taking the derivative with respect to bandwidth $h_j$ gives the statement of the theorem.

**Remark 7.1.** Special treatment is needed if $x$ is a boundary point; see Theorem 2.2 of Ruppert and Wand (1994).

*B. Proof of Theorem 5.2*

Let $\ell$ denote the first row of $S_x$. Then, with $\xi \sim N(0, \sigma^2)$,

$$\begin{aligned}
\widehat{m}_H(x) &= \sum_i \ell_i Y_i = \sum_i \ell_i m(X_i) + \sum_i \ell_i \epsilon_i && (7.46a)\\
&\stackrel{d}{=} \sum_i \ell_i m(X_i) + \xi \sqrt{\sum_i \ell_i^2} && (7.46b)\\
&= \sum_i \ell_i m(X_i) + \frac{\Lambda}{\sqrt{nh_1\cdots h_d}}\,\xi && (7.46c)
\end{aligned}$$

where

$$\Lambda = \sqrt{nh_1\cdots h_d \sum_i \ell_i^2}. \tag{7.47}$$

Thus,

$$\operatorname{Var}(Z_j(t)|X_1,\ldots,X_n) = \sigma^2 \operatorname{Var}\left(\frac{\partial}{\partial h_j}\frac{\Lambda}{\sqrt{nh_1\cdots h_d}}\right). \tag{7.48}$$

Now we find an asymptotic approximation for $\Lambda$.

Recall that

$$S_x = \left(\frac{1}{n}X_x^T W_x X_x\right)^{-1}\frac{1}{n}X_x^T W_x \tag{7.49}$$

and from our previous calculations

$$\Upsilon_n^{-1} = \left(\frac{1}{n}X_x^T W_x X_x\right)^{-1} = \begin{pmatrix} \frac{1}{f(x)} & -\frac{D^T}{f^2(x)} \\ -\frac{D}{f^2(x)} & \frac{H^{-1}}{\nu_2 f(x)} - \frac{DD^T(f(x)-1)}{f^3(x)} \end{pmatrix} + O_P\left((h_j^{(t)})^2/d\right). \tag{7.50}$$

Note that $\sum_i \ell_i^2$ is the (1,1) entry of $S_x S_x^T$. But

$$\begin{aligned} S_x S_x^T &= (\Upsilon^{-1}\frac{1}{n}X_x^T W_x)(\Upsilon^{-1}\frac{1}{n}X_x^T W_x)^T && (7.51a)\\ &= \frac{1}{n^2}\Upsilon^{-1}X_x^T W_x^2 X_x \Upsilon^{-1} && (7.51b)\\ &= \frac{1}{n}\Upsilon^{-1}\begin{pmatrix} \frac{1}{n}\sum_i W_i^2 & \frac{1}{n}\sum_i (X_i-x)^T W_i^2 \\ \frac{1}{n}\sum_i (X_i-x) W_i^2 & \frac{1}{n}\sum_i (X_i-x)(X_i-x)^T W_i^2 \end{pmatrix}\Upsilon^{-1}. && (7.51c)\end{aligned}$$

So $\Lambda^2$ is the $(1,1)$ entry of

$$\begin{aligned} &\Upsilon^{-1}\begin{pmatrix} \frac{h_1\cdots h_d}{n}\sum_i W_i^2 & \frac{h_1\cdots h_d}{n}\sum_i (X_i-x)^T W_i^2 \\ \frac{h_1\cdots h_d}{n}\sum_i (X_i-x) W_i^2 & \frac{h_1\cdots h_d}{n}\sum_i (X_i-x)(X_i-x)^T W_i^2 \end{pmatrix}\Upsilon^{-1} && (7.52a)\\ = \;& \Upsilon^{-1}\begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22}\end{pmatrix}\Upsilon^{-1} && (7.52b)\end{aligned}$$

Next, as in our earlier analysis

$$\begin{aligned} a_{11} &= \int K^2(v)\, f(x - H^{1/2}v)\, dv && (7.53a)\\ &= f(x)\int K^2(v)dv + O_P(h_0^2/d). && (7.53b)\end{aligned}$$

Similarly,

$$a_{21} = \bar{\nu}_2 HD + O_P\left((h_j^{(t)})^4/d)\mathbf{1}\right) \tag{7.54}$$

where $\bar{\nu}_2 I = \int vv^T K^2(v)\, dv$ and

$$a_{22} = f(x)\bar{\nu}_2 H + E \tag{7.55}$$

where $\max_{ij} E_{ij} = O_P\left((h_j^{(t)})^4\right)$. Hence, the leading order expansion of $\Lambda^2$ is given by

$$\frac{\int K^2(v)dv}{f(x)} - \bar{\nu}_2\frac{D^T H D}{f^3(x)} + o(\,\mathrm{tr}(H)) = \frac{\int K^2(v)dv}{f(x)} + O(\,\mathrm{tr}(H)) \tag{7.56}$$

Taking the derivative with respect to $h_j$ we thus conclude that

$$\mathrm{Var}(Z_j(t)|X_1,\ldots,X_n) = \frac{\sigma^2\int K^2(v)\, dv}{f(x)h_j^2}\frac{1}{nh_1\cdots h_d}(1+o_P(1)). \tag{7.57}$$

*C. Proof of Theorem 5.3*

This result characterizes the asymptotic running time, selected bandwidths, and risk of the algorithm. We restate the assumptions made on the functions $m$ and $f$ in Section 5.

(A1) For each $j > r$,

$$\nabla_j \log f(x) = O\left(\frac{1}{n^{1/4}}\right) \tag{7.58}$$

(A2)

$$\liminf_n \min_{1\le j\le r} |m_{jj}(x)| > 0\,. \tag{7.59}$$

To prove the theorem we make use of a version of Mill's inequality, modified for non-zero mean random variables.

**Lemma 7.2.** *Let $Z \sim N(\mu, \sigma^2)$. If $\lambda \ge 2\mu$ and $\lambda^2 \ge 2\sigma^2$ then*

$$\mathbb{P}(|Z| > \lambda) \le \frac{5\lambda}{\sigma} \exp\left\{-\frac{\lambda^2}{8\sigma^2}\right\}. \tag{7.60}$$

*Moreover, if $\lambda \ge 5\sigma$ then*

$$\mathbb{P}(|Z| > \lambda) \le \exp\left\{-\frac{\lambda^2}{16\sigma^2}\right\}. \tag{7.61}$$

*Proof.* Without loss of generality, assume $\mu > 0$. Then,

$$\begin{aligned} \mathbb{P}(|Z| > \lambda) &\le 2\mathbb{P}(Z > \lambda) = 2\mathbb{P}\left(\frac{Z-\mu}{\sigma} > \frac{\lambda-\mu}{\sigma}\right) && \text{(7.62a)} \\ &\le \frac{2\sigma}{\lambda-\mu} \exp\left\{-\frac{(\lambda-\mu)^2}{2\sigma^2}\right\} \equiv B(\mu). && \text{(7.62b)} \end{aligned}$$

Now $B(\mu) = B(0) + \mu B'(\widetilde{\mu})$ for some $0 \le \widetilde{\mu} \le \mu$ and

$$B'(\mu) = \frac{2\sigma}{\lambda-\mu} \exp\left\{-\frac{(\lambda-\mu)^2}{2\sigma^2}\right\} \left(\frac{\lambda-\mu}{\sigma^2} + \frac{1}{\lambda-\mu}\right). \tag{7.63}$$

Hence,

$$B'(\widetilde{\mu}) \le \frac{2\sigma}{\lambda-\mu} \exp\left\{-\frac{(\lambda-\mu)^2}{2\sigma^2}\right\} \left(\frac{\lambda}{\sigma^2} + \frac{1}{\lambda-\mu}\right). \tag{7.64}$$

When $\lambda \ge 2\mu$, $1/(\lambda-\mu) \le 2/\lambda$ and $(\lambda-\mu)^2 \ge \lambda^2/4$ so that if $\lambda^2 \ge 2\sigma^2$ then

$$B'(\widetilde{\mu}) \le \frac{4\sigma}{\lambda} \exp\left\{-\frac{\lambda^2}{8\sigma^2}\right\} \left(\frac{\lambda}{\sigma^2} + \frac{2}{\lambda}\right) \le \frac{8}{\sigma} \exp\left\{-\frac{\lambda^2}{8\sigma^2}\right\}. \tag{7.65}$$

Thus,

$$\mathbb{P}(|Z| > \lambda) \le \frac{2\sigma}{\lambda} \exp\left\{-\frac{\lambda^2}{2\sigma^2}\right\} + \frac{8\mu}{\sigma} \exp\left\{-\frac{\lambda^2}{8\sigma^2}\right\} \le \frac{5\lambda}{\sigma} \exp\left\{-\frac{\lambda^2}{8\sigma^2}\right\}. \tag{7.66}$$

The last statement follows since $5xe^{-x^2/8} \le e^{-x^2/16}$ for all $x \ge 5$. $\square$

*Proof* of Theorem 5.3. First consider $j > r$. Let $V_t = \{j > r : \ h_j = h_0\beta^t\}$ be the set of irrelevant dimensions that are active at stage $t > 1$ of the algorithm. From Assumption (A1) it follows that, for sufficiently large $n$,

$$\lambda_j \ge |L_j^{(t)}| = O(h_0^3 n^{-1/2}) \tag{7.67}$$

for $j > r$. Also, $\lambda_j \ge 5s_j^{(t)}$. From Theorems 5.1 and 5.2, we have that, with probability tending to one uniformly over $j > r$, $\lambda_j \ge 2|\mu_j^{(t)}|$ and $\lambda_j^2 \ge 2v_j^{(t)}$, and $v_j^{(t)}/(s_j^{(t)})^2 = 1 + o(1)$. Thus, by Lemma 7.2,

$$\begin{aligned}
\mathbb{P}(|Z_j| > \lambda_j, \text{ for some } j \in V_t) &\le \sum_{j \in V_t} \mathbb{P}(|Z_j| > \lambda_j) + o(1) && (7.68a)\\
&\le d \exp\left(-\lambda_j^2/(16 v_j)\right) + o(1) && (7.68b)\\
&= d \exp\left(-\lambda_j^2(1 + o(1))/(16 s_j^2)\right) + o(1) \to 0. && (7.68c)
\end{aligned}$$

Therefore, with probability tending to one, $h_j = h_0$ for each $j > r$, meaning that the bandwidth for each irrelevant dimension is frozen in the first step in the algorithm.

Now consider $j \le r$. By (5.12), and Theorem 5.1, for all large $n$,

$$|\mu_j| \ge |L_j^{(t)}| - |\mu_j - L_j^{(t)}| = |L_j^{(t)}| - O_P((h_j^{(t)})^4/d) \ge c h_j^{(t)} |m_{jj}(x)| \tag{7.69}$$

for some $c > 0$. Without loss of generality, assume that $c h_j m_{jj}(x) > 0$. We claim that in iteration $t$ of the algorithm, if

$$t \ \le \ \frac{1}{4 + r} \log_{1/\beta}\left(\frac{c^2 n A_{\min}^2 h_0^{d+4}}{8C \log(n c_n)}\right) = t_0 \tag{7.70}$$

then

$$\mathbb{P}(h_j = h_0\beta^t, \text{ for all } j \le r) \longrightarrow 1. \tag{7.71}$$

To show this, first note that (7.70) can be written as

$$\left(\frac{1}{\beta}\right)^{t(4+r)} \ \le \ \frac{c^2 n A_{\min}^2 h_0^{d+4}}{8C \log(c_n n)}. \tag{7.72}$$

Except on an event of vanishing probability, we have shown above that

$$\prod_{j>r} \frac{1}{h_j} \ = \ \frac{1}{h_0^{(d-r)}}. \tag{7.73}$$

So on the complement of this event, if each relevant dimension is active at step $s \leq t$, we have

$$\begin{aligned}
\frac{\lambda_j^2}{h_j^2} &= \frac{2s_j^2 \log(nc_n)}{h_j^2} = \frac{2C\log(nc_n)}{nh_j^4} \prod_i \frac{1}{h_i} && (7.74a)\\
&\leq \frac{2C\log(nc_n)}{nh_0^{d+4}} \left(\frac{1}{\beta}\right)^{(4+r)t} \leq \frac{c^2 A_{\min}^2}{4} && (7.74b)\\
&\leq \frac{c^2 m_{jj}(x)^2}{4} && (7.74c)
\end{aligned}$$

which implies that

$$cm_{jj}(x)h_j \geq 2\lambda_j \tag{7.75}$$

and hence

$$\frac{cm_{jj}(x)h_j - \lambda_j}{s_j} \geq \frac{\lambda_j}{s_j} = \sqrt{2\log(nc_n)} \tag{7.76}$$

for each $j \leq r$. Now,

$$\begin{aligned}
\mathbb{P}(\text{rodeo halts}) &= \mathbb{P}(|Z_j| < \lambda_j \text{ for all } j \leq r) \leq \mathbb{P}(|Z_j| < \lambda_j \text{ for some } j \leq r) && (7.77a)\\
&\leq \sum_{j\leq r} \mathbb{P}(|Z_j| < \lambda_j) \leq \sum_{j \leq r} \mathbb{P}(Z_j < \lambda_j) && (7.77b)\\
&\leq \sum_{j\leq r} \mathbb{P}\left(\frac{Z_j - \mu_j}{s_j} > \frac{\mu_j - \lambda_j}{s_j}\right) && (7.77c)\\
&\leq \sum_{j\leq r} \mathbb{P}\left(\frac{Z_j - \mu_j}{s_j} > \frac{cm_{jj}(x)h_j - \lambda_j}{s_j}\right) && (7.77d)\\
&\leq \frac{r}{nc_n\sqrt{2\log(nc_n)}}. && (7.77e)
\end{aligned}$$

Finally, summing over all iterations $s \leq t$ gives

$$\begin{aligned}
\mathbb{P}\left(\bigcup_{s\leq t}\bigcup_{j\leq r}\left\{|Z_j^{(t)}| < \lambda_j^{(t)}\right\}\right) &\leq \frac{tr}{nc_n\sqrt{2\log(nc_n)}} && (7.78a)\\
&\leq \frac{\frac{r}{r+4}\log_{1/\beta}\left(\frac{c^2 nA_{\min}^2 h_0^{d+4}}{8C\log(nc_n)}\right)}{nc_n\sqrt{2\log(nc_n)}} \longrightarrow 0. && (7.78b)
\end{aligned}$$

Thus, the bandwidths $h_j$ for $j \leq r$ satisfy, with high probability,

$$\begin{aligned}
h_j = h_0\beta^t &\leq h_0\beta^{t_0} && (7.79a)\\
&= h_0\left(\frac{8C\log(nc_n)}{c^2 A_{\min}^2 nh_0^{d+4}}\right)^{1/(4+r)} && (7.79b)\\
&= n^{-1/(4+r)}\left(\frac{8C\log(nc_n)}{c^2 A_{\min}^2 h_0^{d-r}}\right)^{1/(4+r)}. && (7.79c)
\end{aligned}$$

Now, since $d \log d = O(\log n)$ it follows that $d = O(\log n / \log\log n)$ and

$$\begin{aligned} \frac{1}{h_0^d} &= O\left((\log\log n)^d\right) && (7.80a) \\ &= O\left(n^{\frac{\log\log\log n}{\log\log n}}\right) && (7.80b) \\ &= O\left(n^{\epsilon}\right) && (7.80c) \end{aligned}$$

for any $\epsilon > 0$. In particular, with probability approaching one, the algorithm runs for a number of iterations $T_n$ bounded from below by

$$T_n \geq \frac{1}{4+r} \log_{1/\beta}(n^{1-\epsilon} a_n) \tag{7.81}$$

where

$$a_n \;=\; \frac{c^2 A_{\min}^2}{8C \log(n c_n)} \;=\; \widetilde{\Omega}(1)\,. \tag{7.82}$$

We next show that the algorithm is unlikely to reach iteration $s$, if

$$s \;\geq\; \frac{1}{4+r} \log_{1/\beta}\left(n^{1+\delta} b_n\right) \tag{7.83}$$

for any $\delta > 0$, where $b_n = \widetilde{O}(1)$ is defined below. From the argument above, we know that except on an event of vanishing probability, each relevant dimension $j \leq r$ has bandwidth no larger than

$$\begin{aligned} h_j &\leq h_0 \beta^{\frac{1}{4+r}} \log_{1/\beta}(n^{1-\epsilon} a_n) && (7.84a) \\ &= \frac{h_0}{(n^{1-\epsilon} a_n)^{1/(4+r)}}. && (7.84b) \end{aligned}$$

Thus, if relevant dimension $j$ has bandwidth $h_j \leq h_0 \beta^s$, then from Theorem 5.2 we have that

$$\begin{aligned} \frac{s_j^2}{A_j^2 h_j^2} &\geq \frac{C}{A_j^2 n h_0^4 \beta^{4s}} \frac{n^{r(1-\epsilon)/(4+r)} a_n^{r/(4+r)}}{h_0^r} \frac{1}{h_0^{d-r}} && (7.85a) \\ &= \frac{C}{A_j^2 n^{4/(4+r)+\delta}} \frac{a_n^{r/(4+r)}}{h_0^{4+d}} \frac{1}{\beta^{4s}} && (7.85b) \\ &\geq \frac{C}{A_{\max}^2 n^{4/(4+r)+\delta}} \frac{a_n^{r/(4+r)}}{h_0^{4+d}} \frac{1}{\beta^{4s}} && (7.85c) \end{aligned}$$

where $\delta = r\epsilon/(4+r)$. Therefore,

$$\frac{s_j^2}{A_j^2 h_j^2} \geq \log\log n \tag{7.86}$$

in case

$$\left(\frac{1}{\beta}\right)^s \geq (n^{1+\delta}b_n)^{1/(4+r)} \tag{7.87a}$$

$$= n^{1/(4+r)+\delta'}\left(\frac{A_{\max}^2 \log\log n}{C a_n^{r/(4+r)}}\right)^{1/4} \tag{7.87b}$$

$$\geq n^{1/(4+r)+\delta'}\left(\frac{A_{\max}^2 h_0^{4+d} \log\log n}{C a_n^{r/(4+r)}}\right)^{1/4} \tag{7.87c}$$

which defines $b_n = \widetilde{O}(1)$. It follows that in iteration $s \geq \frac{1}{4+r}\log_{1/\beta}\left(n^{1+\delta}b_n\right)$, the probability of a relevant variable having estimated derivative $Z_j$ above threshold is bounded by

$$\mathbb{P}(|Z_j| > \lambda_j, \text{ for some } j \leq r) \leq \sum_{j\leq r}\mathbb{P}\left(\frac{|Z_j|}{s_j} > \frac{\lambda_j}{s_j}\right) \tag{7.88a}$$

$$\leq \sum_{j\leq r}\left(\frac{s_j}{\lambda_j}e^{-\lambda_j^2/(2s_j^2)} + \frac{1}{4}\frac{\mu_j^2}{s_j^2}\right) \tag{7.88b}$$

$$\leq \frac{r}{nc_n\sqrt{2\log(nc_n)}} + \frac{1}{4}\sum_{j\leq r}\frac{\mu_j^2}{s_j^2} \tag{7.88c}$$

$$\leq \frac{r}{nc_n\sqrt{2\log(nc_n)}} + \frac{r}{4\log\log n} \tag{7.88d}$$

$$= O\left(\frac{1}{\log\log n}\right) \tag{7.88e}$$

which gives the statement of the theorem. □

*Proof* of Corollary 5.4. Given the bandwidths in (5.15) and (5.16) we have that the squared (conditional) bias is given by

$$\text{Bias}^2(\widehat{m}_{h^\star}) = \left(\sum_{j\leq r} A_j h_j^{\star 2}\right)^2 + o_P(\text{tr}(H^\star)) \tag{7.89a}$$

$$= \sum_{i,j\leq r} A_i A_j h_i^{\star 2} h_j^{\star 2} + o_P(\text{tr}(H^\star)) \tag{7.89b}$$

$$= O_P(n^{-4/(4+r)+\epsilon}) \tag{7.89c}$$

by Theorem 5.3. Similarly, from the proof of Theorem 5.2 the (conditional) variance is

$$\mathbb{V}\text{ar}(\widehat{m}_{h^\star}) = \frac{1}{n}\left(\prod_i \frac{1}{h_i^\star}\right)\frac{R(K)}{f(x)}\sigma^2(1+o_P(1)) \tag{7.90a}$$

$$= O_P(n^{-1+r/(r+4)+\epsilon}) \tag{7.90b}$$

$$= O_P(n^{-4/(4+r)+\epsilon}) \tag{7.90c}$$

where $R(K) = \int K(u)^2\,du$. The result follows from the bias-variance decomposition. $\square$

*D. Proof of Theorem 5.5*

This result shows that assumption (A1) can be dropped if we modify the algorithm by simply increasing the threshold slightly. Examples of choices for $\rho_n$ are $\rho_n = \delta h_0^2$ or $\rho_n = h_0^{3-\delta}$ for some small $\delta > 0$.

Let $j > r$. Without loss of generality assume that $\mu_j \geq 0$. Recall that, if variable $j$ is still active in round $t$, $\mu_j = ch_j^3 + O_P((h_j^{(t)})^2/d)$ for some constant $c > 0$. It follows from (5.20) that, for all large $n$,

$$\lambda_j > \mu_j + s_j\sqrt{2\log(nc_n)} \tag{7.91}$$

and hence

$$\frac{s_j}{\lambda_j - \mu_j} e^{-(\lambda_j-\mu_j)^2/(2s_j^2)} \leq \frac{1}{nc_n\sqrt{2\log(nc_n)}}. \tag{7.92}$$

Then, with $\xi \sim N(0,1)$, we have

$$\begin{aligned}
\mathbb{P}(|Z_j| > \lambda_j, \text{ for some j} > \text{r}) &\leq \sum_j 2\mathbb{P}\left(\xi > \frac{\lambda_j - \mu_j}{s_j}\right) && (7.93a)\\
&\leq \sum_j \frac{2s_j}{\lambda_j - \mu_j} e^{-(\lambda_j-\mu_j)^2/(2s_j^2)} && (7.93b)\\
&\leq \frac{2d}{nc_n\sqrt{2\log(nc_n)}} \longrightarrow 0. && (7.93c)
\end{aligned}$$

Now consider $j \leq r$. It follows from (5.20) that, for all large $n$,

$$\frac{\mu_j - \rho_n\beta^{3t}}{s_j} > c\frac{\mu_j}{s_j} \tag{7.94}$$

for some $c > 0$. Hence,

$$\begin{aligned}
\mathbb{P}(Z_j < \lambda_j) &= \mathbb{P}\left(\frac{Z_j - \mu_j}{s_j} < \frac{\lambda_j - \mu_j}{s_j}\right) && (7.95a)\\
&= \mathbb{P}\left(\xi > \frac{\mu_j - \lambda_j}{s_j}\right) && (7.95b)\\
&\leq \mathbb{P}\left(\xi > \frac{c\mu_j - \lambda_j'}{s_j}\right) && (7.95c)
\end{aligned}$$

where $\lambda_j' = s_j\sqrt{2\log(nc_n)}$. The calculation then proceeds as before.

## VIII. Acknowledgments

We thank Gilles Blanchard and the editors and referees for many helpful comments on this work.

## References

Breiman, L., Friedman, J. H., Olshen, R. A. and Stone, C. J. (1984). *Classification and regression trees*. Wadsworth Publishing Co Inc.

Bühlmann, P. and Yu, B. (2005). Boosting, model selection, lasso and nonnegative garrote. Tech. rep., Berkeley.

Donoho, D. (2004). For most large underdetermined systems of equations, the minimal $\ell^1$-norm near-solution approximates the sparest near-solution. Tech. rep., Stanford University.

Efron, B., Hastie, T., Johnstone, I. and Tibshirani, R. (2004). Least angle regression. *The Annals of Statistics* **32** 407–499.

Fan, J. (1992). Design-adaptive nonparametric regression. *Journal of the American Statistical Association* **87** 998–1004.

Fan, J. and Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association* **96** 1348–1360.

Fan, J. and Peng, H. (2004). Nonconcave penalized likelihood with a diverging number of parameters. *The Annals of Statistics* **32** 928–961.

Friedman, J. H. (1991). Multivariate adaptive regression splines. *The Annals of Statistics* **19** 1–67.

Fu, W. and Knight, K. (2000). Asymptotics for lasso type estimators. *The Annals of Statistics* **28** 1356–1378.

George, E. I. and McCulloch, R. E. (1997). Approaches for Bayesian variable selection. *Statistical Sinica* **7** 339–373.

Girosi, F. (1997). An equivalence between sparse approximation and support vector machines. Tech. rep., Massachusetts Institute of Technology. A.I. Memo No. 1606.

Györfi, L., Kohler, M., Krzyżak, A. and Walk, H. (2002). *A Distribution-Free Theory of Nonparametric Regression*. Springer-Verlag.

Hastie, T. and Loader, C. (1993). Local regression: Automatic kernel carpentry. *Statistical Science* **8** 120–129.

Hastie, T., Tibshirani, R. and Friedman, J. H. (2001). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer-Verlag.

Hristache, M., Juditsky, A., Polzehl, J. and Spokoiny, V. (2001). Structure adaptive approach for dimension reduction. *Ann. Statist.* **29** 1537–1566.

Lawrence, N. D., Seeger, M. and Herbrich, R. (2003). Fast sparse Gaussian process methods: The informative vector machine. In *Advances in Neural Information Processing Systems 15*. MIT Press.

Lepski, O. V., Mammen, E. and Spokoiny, V. G. (1997). Optimal spatial adaptation to inhomogeneous smoothness: An approach based on kernel estimates with variable bandwidth selectors. *The Annals of Statistics* **25** 929–947.

Li, L., Cook, R. D. and Nachsteim, C. (2005). Model-free variable selection. *J. R. Statist. Soc. B.* **67** 285–299.

Rice, J. (1984). Bandwidth choice for nonparametric regression. *The Annals of Statistics* **12** 1215–1230.

Ruppert, D. (1997). Empirical-bias bandwidths for local polynomial nonparametric regression and density estimation. *Journal of the American Statistical Association* **92** 1049–1062.

Ruppert, D. and Wand, M. P. (1994). Multivariate locally weighted least squares regression. *The Annals of Statistics* **22** 1346–1370.

Samarov, A., Spokoiny, V. and Vial, C. (2005). Component identification and estimation in nonlinear high-dimensional regression models by structural adaptation. *J. Amer. Statist. Assoc.* **100** 429–445.

Smola, A. and Bartlett, P. (2001). Sparse greedy Gaussian process regression. In *Advances in Neural Information Processing Systems 13*. MIT Press.

Stone, C. J., Hansen, M. H., Kooperberg, C. and Truong, Y. K. (1997). Polynomial splines and their tensor products in extended linear modeling (Disc: p1425-1470). *The Annals of Statistics* **25** 1371–1425.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B, Methodological* **58** 267–288.

Tipping, M. (2001). Sparse Bayesian learning and the relevance vector machine. *Journal of Machine Learning Research* **1** 211–244.

Tropp, J. A. (2004). Greed is good: Algorithmic results for sparse approximation. *IEEE Trans. Inform. Theory* **50** 2231–2241.

Tropp, J. A. (2006). Just relax: Convex programming methods for identifying sparse signals. *IEEE Trans. Inform. Theory* **51** 1030–1051.

Turlach, B. (2004). Discussion of "least angle regression". *The Annals of Statistics* **32** 494–499.

Zhang, H., Wahba, G., Lin, Y., Voelker, M., Ferris, R. K. and Klein, B. (2005). Variable selection and model building via likelihood basis pursuit. *J. of the Amer. Stat. Assoc.* **99** 659–672.